

# Persuading Investors: A Video-Based Study\*

ALLEN HU                      and                      SONG MA

Persuasive communication functions through not only content but also delivery, i.e., facial expression, tone of voice, and diction. This paper examines the persuasiveness of delivery in startup pitches. Using machine learning algorithms to process full pitch videos, we quantify persuasion in visual, vocal, and verbal dimensions. We find that positive (i.e., passionate, warm) pitches increase funding probability. Yet conditional on funding, startups with higher levels of pitch positivity underperform. Women are more heavily judged on delivery when evaluated in single-gender teams, but they are neglected when co-pitching in mixed-gender teams. Using an experiment, we show that persuasion delivery works mainly through leading investors to form inaccurate beliefs.

**JEL Classification:** D91, G41, C55, G24

**Keywords:** Persuasion, Behavioral Economics, Video Data, Machine Learning

---

\*Hu is with Sauder School of Business, University of British Columbia. Ma is with Yale University and NBER. We thank Wei Xiong (the Editor), an anonymous Associate Editor, and two anonymous referees. We also thank seminar participants at ASSA, Boston College, Chicago Booth, Colorado, EFA, Finance in the Cloud, Georgia State, IÉSEG, Indiana, MFA, NBER Behavioral Finance Meeting, NBER Summer Institute (Corporate Finance; Entrepreneurship), RFS/GSU FinTech Conference, SFS Cavalcade, Washington St. Louis, WFA, and Yale and numerous colleagues for helpful comments. Yajie Luo and Ran You provided excellent research assistance. The project received IRB approval at Yale University. All errors are our own. We have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose. The paper was previously circulated under the title “Human Interactions and Financial Investment: A Video-Based Approach.” We have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Correspondence: Song Ma, Yale School of Management, 165 Whitney Avenue, New Haven, CT 06511, USA; e-mail: song.ma@yale.edu.

Many decisions in firms and financial markets are made after interpersonal persuasion, for example, pitches to investors, board presentations, and job interviews. It is often believed that in such persuasive communications, investors or interviewers not only learn information from the content (e.g., the numbers, financial models, market analysis)<sup>1</sup> but are also swayed by the rich characteristics of how the persuasion is delivered—features like facial expression, tone of voice, or diction of speech can be impactful. As Mehrabian (1971)’s study shows, there is a 7-55-38 rule of communication: words convey 7 percent of a message, body language such as facial expressions accounts for 55 percent, and tone delivers 38 percent.<sup>2</sup>

Research on the effect of non-verbal features in communications has recently started to emerge in financial economics. Mayew and Venkatachalam (2012) and Gorodnichenko, Pham, and Talavera (2023) study vocal tones of CEOs in earnings conferences and that of Federal Reserve chairs in press conferences, respectively, and show that the market extracts information from vocal cues. Some studies, however, cast doubt on the value of using these communication features in economic decision-making. For example, using executive assessment data, Kaplan, Klebanov, and Sorensen (2012) and Kaplan and Sorensen (2021) show that boards overweight interpersonal communication skills in their executive hiring decisions at the cost of poor future managerial performance. Open questions remain concerning the role of persuasive communication features in financial decision-making, accounting for comprehensive dimensions like facial expressions, vocal tones, and verbal cues; whether these features lead to better or biased investment decisions; and the underlying mechanisms.

In this paper, we make progress in answering these questions in three dimensions. First, we introduce a method to process video data of persuasive communications, from which we can simultaneously quantify multi-modal characteristics including visual (e.g., facial expressions), vocal (e.g., tone of voice), and verbal (e.g., word choices in the script) features. Second, we use a non-lab, real-world setting of interpersonal persuasion in which entrepreneurs pitch to experienced investors as a key step for obtaining early-stage funding. This setting allows us to closely link features in the

---

<sup>1</sup>DellaVigna and Gentzkow (2010) survey persuasion models that mostly focus on how the content in persuasive interactions matters for decisions. The content may be informational, such as a project’s net present value (Stigler, 1961; Milgrom and Roberts, 1986; Dewatripont and Tirole, 1999; Kamenica and Gentzkow, 2011). Conversely, the content may be noninformational, such as “framing” the pitch (Mullainathan, Schwartzstein, and Shleifer (2008)), using appealing peripheral content that caters to people’s intuition and attracts attention (Bertrand et al. (2010)), or using “models” that lead receivers to interpret data and facts in a certain way (Schwartzstein and Sunderam (2021)).

<sup>2</sup>The debate on exact percentages for each channel is continuing, and these percentages could be different setting by setting. But it is clear that effective persuasion relies on more than just hard information conveyed in words.

pitch with venture investors’ investment decisions and to track the long-term performance of the companies, if funded, whose founders pitch differently. Third, we conduct an experiment on MBA students to examine potential mechanisms through which features in the pitches swing investor decisions, in particular, through changing investors’ beliefs or catering to investors’ preferences.

Our analysis employs a setting in which early-stage startups pitch to large and highly ranked venture accelerators in the US. This setting is important given the value of entrepreneurship and innovation for the economy (Kortum and Lerner, 2000; Haltiwanger et al., 2013) and our general lack of understanding on how the venture selection process works (Bernstein, Korteweg, and Laws, 2017; Ewens and Townsend, 2020). We collect more than 1,000 pitch videos when these startups sought funding from an investor. These are augmented by detailed characteristics on the startups (e.g., industry, location) and on the founders (e.g., work experience, education). We also track investor decisions subsequent to the pitches and future startup performance.

To quantify pitch delivery, we exploit a set of machine learning (ML) algorithms to simultaneously quantify features along visual, vocal, and verbal dimensions using video data. This method is efficient, reproducible, and easy to be adapted for future research. It proceeds in three steps.<sup>3</sup> First, we project pitch videos onto what we call the “three-V” spaces: visual, vocal, and verbal. Given the purpose of capturing the whole persuasion process, we process the full video of each pitch. The visual dimension is represented as a series of images at a frequency of ten frames per second (i.e., 600 images for a one-minute pitch). We extract the full audio stream from the video. The audio file encompasses both the vocal component and the verbal script generated through a speech-to-text algorithm.

Next, we construct persuasion delivery features from these three-V channels using ML algorithms. We use easily accessible computation services, such as Face++, Microsoft Azure, and Google Cloud, to perform the computation and construct measures. The key algorithms are trained, tested, and cross-checked by reputable providers using millions of human-rated training observations. Using these algorithms also allows for high replicability and transparency, and it lowers the computation burdens for researchers. These algorithms generate detailed pitch features including visual emotions (e.g., positive, negative), vocal emotions (e.g., positive, negative, valence, arousal), textual sentiment (e.g., positive, negative), and psychological features (e.g., warmth).

In the last step, applying factor analysis to all the detailed features, we create an overall measure,

---

<sup>3</sup>A complete coding example following these steps is provided in Appendix B.

the Pitch Factor, that summarizes “how well” the startup team delivers the pitch. This is similar to that in [Tetlock \(2007\)](#), [Kaplan and Sorensen \(2021\)](#), and [Kaplan, Klebanov, and Sorensen \(2012\)](#) and is common in unstructured data analysis. This process collapses detailed features into a single factor that captures the maximum variance in the set of pitch features. Empirically, the Pitch Factor loads positively on positivity dimensions and negatively on negativity dimensions, so, intuitively, the measure can be interpreted as the overall level of positivity—for example, happiness, passion, warmth, enthusiasm—in a pitch.

The Pitch Factor, though being created using ML and statistical methods, corresponds well with the entrepreneurial pitch setting and finds its roots in prior economic research. In a large-scale VC survey conducted by [Gompers, Gornall, Kaplan, and Strebulaev \(2020\)](#), passion, broadly defined, is consistently ranked a top-three factor when selecting portfolio companies. Indeed, under the general impression that entrepreneurs are positive, energetic, and optimistic ([Åstebro, Herz, Nanda, and Weber \(2014\)](#)), the positivity feature revealed in the pitch may have the power to swing the opinions of agents thinking categorically and coarsely ([Fryer and Jackson, 2008](#); [Mullainathan et al., 2008](#)). Additionally, positivity demonstrated in pitches may be contagious and particularly salient in affecting investors’ emotional states ([DellaVigna \(2009\)](#)), which in turn influences both the beliefs of future prospects and risk assessments.<sup>4</sup>

We uncover four main findings. First, startup teams that score a higher Pitch Factor—that is, those showing more positivity, passion, and enthusiasm in their pitches—are more likely to obtain funding. This pattern is consistent across individual measures from the vocal, visual, and verbal dimensions. A one-standard-deviation change in the more passionate direction is associated with a three percentage point increase in the probability of receiving funding, or a 35.2 percent increase from the baseline probability of receiving funding. The results are robust to controlling for a set of text-based measures of pitch content on idea novelty and linguistic features. The results also do not appear to be driven by the sample selection issues concerning the availability of videos in data construction.

The analysis also demonstrates the advantage of using the video data at full length and of simultaneously considering the three-V dimensions, a key difference between our approach with

---

<sup>4</sup>For emotional contagion, see [Hatfield, Cacioppo, and Rapson \(1993\)](#) and [Barsade \(2002\)](#). For emotions, beliefs, and their impact in economics, see [Johnson and Tversky \(1983\)](#), [Arkes, Herren, and Isen \(1988\)](#), [Clare et al. \(1994\)](#), [Loewenstein et al. \(2001\)](#), [Hirshleifer and Shumway \(2003\)](#), [Dahl and DellaVigna \(2009\)](#), and [Kuhnen and Knutson \(2011\)](#).

prior work. When running horse-race models, the Pitch Factor generated using the full video dwarfs that are constructed using thin slices of videos. Similarly, measures constructed from individual V-channels are less robust compared to the Pitch Factor. We also show that the full-video Pitch Factor explains the largest amount of variations in investors' funding decisions when being compared with thin-sliced or single-channel factors.

The impact of pitch delivery features does not seem to be simply explained by the omission of unobserved startup and founder controls in the above baseline analysis. The concern is that if the style of pitches correlates with entrepreneur-level quality traits (say, better founders communicate more passionately), the pitch feature may simply be picking up omitted quality metrics. We adopt a test following [Altonji, Elder, and Taber \(2005\)](#) and [Oster \(2019\)](#) and include an extensive set of controls used in the literature, such as founders' education, employment background, and startup experience, in the investment regression ([Bernstein, Korteweg, and Laws, 2017](#)). When we do so, the estimated impact of pitch features remains stable in economic magnitude and statistical significance. The statistical tests show that the impact of pitch features is robust to a wide set of reasonable parameters regarding omitted startup quality.

Second, we examine whether the first “better pitch, more funding” finding is driven by investors incorporating pitch delivery features to improve their investment decisions—and we find little support for this view. If the above finding simply reflects that a person who makes a better pitch runs a better startup, then the invested companies with higher levels of pitch positivity would likely perform better than those with poorer pitch features ([Fisman, Paravisini, and Vig, 2017](#); [Ewens and Townsend, 2020](#)). To examine this, we use the long-term performance of startups and track their performance using employment, development, and survival (tracked through startup website update activities), the probability of attracting follow-up financing and the long-term funding amount, and the probability of achieving a milestone exit like an IPO or an acquisition.

We find that none of the positive pitch features link to better long-term performance. In fact, many positive pitch characteristics are associated with poorer long-term performance. This analysis certainly does not intend to establish any causal interpretation between pitch features and performance. Instead, our preferred interpretation is that investors lower the bar (or, equivalently, assign a high investment probability) for startups that show more positivity and passion in their pitch delivery, which lowers the portfolio's true average project success probability.

Third, we provide evidence that gender plays a role in how the Pitch Factor impacts funding,

which is in a direction consistent with gender biases. Understanding gender gaps in entrepreneurship is not only socially important (Ewens (2022)) but also helpful for us to further isolate the potential biases when incorporating pitch delivery in making investment decisions. Previous research shows that women are more often judged based on their appearance and not on substance (Fredrickson and Roberts (1997)). In addition, women and men are expected to follow different gender stereotypes (Bordalo et al. (2019))—for example, women are expected to portray warmth, empathy, and altruism more than men are (Kite, Deaux, and Haines, 2008; Ellemers, 2018). Women also receive less recognition in group work (Sarsons et al. (2021)). Biasi and Sarsons (2022) and Cullen and Perez-Truglia (2023) show that social interactions with managers often put female employees in a disadvantageous position compared to men. In entrepreneurial investment settings, Ewens and Townsend (2020) and Hebert (2020) show that venture investors have gender stereotypes that can cause them to hold some bias against women.

We show that when single-gender teams pitch to investors, women and men are both evaluated on the Pitch Factor, but with different intensities. The penalty (i.e., a decrease in funding probability) for being one standard deviation less passionate and positive is nine times bigger for women than for men. This result does not seem to be explained away by different speaking styles, industry compositions of startups, or algorithmic accuracy across genders. Next, we find that in mixed-gender teams, men’s pitch features remain relevant, but women’s pitch features become statistically and economically irrelevant. This suggests that women are essentially overlooked when co-presenting with their male teammates.

Finally, we conduct an experiment to investigate the economic mechanisms through which persuasion delivery features influence investors’ decisions. We follow DellaVigna and Gentzkow (2010) and broadly categorize the potential mechanisms to (inaccurate) belief-based and preference/taste-based. We try to separate them using a venture investing experiment designed following the structure of Bohren, Haggag, Imas, and Pope (2019). In the experiment, we ask MBA students to watch ten pitch videos and make investment decisions to maximize their payoff. Importantly, we also directly elicit their beliefs on the success probability of startups that are pitching. This experimental design allows us to separate the potential mechanisms using a mediation analysis framework (VanderWeele (2016)).

Consistent with the inaccurate beliefs channel, we find that investors mistakenly think that startups with more positive pitch features are more likely to succeed even though the realized

performance of those companies is not higher, as discussed above—hence the inaccurate beliefs. After controlling for the elicited belief that the investment decision is based on, pitch features remain influential as a standalone determinant, consistent with the preference-based channel. Our quantitative decomposition shows that the inaccurate beliefs and preference-based channels contribute 80 percent and 20 percent, respectively, to the relation between the non-content persuasion features and investment decisions.

The contribution of this paper is twofold. First, for entrepreneurial finance, we leverage ML techniques, video data, and experiments to provide novel evidence on how non-verbal communication features affect early-stage venture investment decisions, how they lead to biases and underperformance, and importantly, the underlying economic mechanisms. Second, in terms of method, we provide, to our knowledge, the first comprehensive application of video-processing techniques in economic research that can simultaneously incorporate visual, vocal, and verbal information and can incorporate full-length video. This technique can open valuable research venues for researchers exploring other questions.

This paper relates to the literature in entrepreneurial finance on how VC investors make decisions. [Gompers et al. \(2020\)](#) conduct a recent large-scale survey to study VCs’ decision-making and find that “passion” is a top consideration when evaluating startup teams. [Bernstein et al. \(2017\)](#), using a randomized experiment, also find that hard information about founding teams is key for early-stage investors. [Kaplan and Strömberg \(2004\)](#) identify a broader set of startup characteristics that are considered by VCs. Our paper adds to the literature by focusing specifically on the pitch, which is an important step in the decision process but was challenging to study in the past.

Our results suggest a potential bias arising from interpersonal communication with teams, or soft information in general. This is closely related to recent work documenting biases arising from evaluating startup teams. [Kaplan et al. \(2009\)](#) study the life cycle of startups and argue that, at the margin, investors should place more weight on the business. [Ewens and Townsend \(2020\)](#) and [Hebert \(2020\)](#) document the potential gender bias when venture investors make decisions. Using executive assessment data, [Kaplan and Sorensen \(2021\)](#) show that boards overweight interpersonal skills in their CEO hiring decisions—interpersonal skills lead to higher hiring probability but are negatively related to subsequent performance in VC-funded firms. In a recent paper [Huang et al. \(2023\)](#) use a survey-based method to capture static first-impression features developed from video stills (pictures) on the entrepreneurial TV shows Shark Tank and Startup Battlefield. Complementary



to our work, the authors find that first impressions of entrepreneurs’ facial traits are associated with favorable decisions from judges on the TV show. In contrast, our ML-based approach examines dynamic communication features across all three channels—visual, vocal, and verbal—controls for the pitch’s rich informational content, and analyzes a real-world investment setting rather than (partially) scripted TV shows. Our results also suggest that analyzing all three channels is crucial for capturing the full range of communication features and studying single channels is likely to produce incomplete results. Additionally, unlike [Huang et al. \(2023\)](#), our experiment analysis also sheds light on the economic mechanism (inaccurate belief formation) behind the findings.

Method-wise, a key contribution of our approach is the ability to jointly use information from all three information channels. Most of this work focuses on single-channel analysis. [Gentzkow, Kelly, and Taddy \(2019\)](#) thoroughly review the progress in using textual data in economics and finance in the past decade.<sup>5</sup> [Hobson, Mayew, and Venkatachalam \(2012\)](#), [Mayew and Venkatachalam \(2012\)](#), and [Gorodnichenko, Pham, and Talavera \(2023\)](#) study vocal cues. A nascent literature has started to use ML-based algorithms to code static facial traits from images ([Boxell, 2024](#); [Joo and Steinert-Threlkeld, 2018](#); [Peng, 2018](#); [Choudhury et al., 2019](#); [Peng et al., 2021](#)), say “attractiveness” and “dominance” rather than facial movements and expressions.

Using three-V channels to construct the Pitch Factor shows its empirical and economic value. Empirically, we find that using three-V information is more robust and accounts for a larger variation in funding decisions than measures constructed from individual V channels. Employing three-V channels also provides a key economic insight compared to the single-channel analysis: when studying the influence of communication features, it is important to consider not only the three-V channels individually but also the covariance matrix of information in these channels. A natural implication is that our finding is less likely to be confounded by omitted variable bias compared to studies using a single dimension alone. For example, if positive facial expressions positively correlate with other delivery features (e.g., passionate voices), then correctly estimating the effect of facial expressions becomes difficult when only analyzing the facial information without accounting for other channels.

A separate stream of work, mainly in behavioral sciences and political science, has shown the usefulness of videos in empirical research. Instead of leveraging data science techniques and

---

<sup>5</sup>For textual analysis in specific research settings, see [Antweiler and Frank \(2004\)](#), [Tetlock \(2007\)](#), [Tetlock et al. \(2008\)](#), [Hoberg and Phillips \(2010\)](#), [Loughran and McDonald \(2011\)](#), [Hoberg and Phillips \(2016\)](#), and [Loughran and McDonald \(2016\)](#), among others.



computation, these studies generally use the “subject rating, thin-sliced data, static perception” norm. That is, researchers recruit subjects to view thin-sliced data, most often still images (thus no movements or audio information), and sometimes also very short video clips. The subjects are then asked to rate static features of the speaker such as attractiveness, trustworthiness, and competence; these features in turn are correlated with outcomes.<sup>6</sup>

By introducing ML techniques, our method improves upon two dimensions. First, it goes beyond rating static perceptions and quantifies the complete persuasion process based on complete vocal information and facial movements. Roughly speaking, given a set of static personal features (e.g., one with an above-average appearance), our method focuses on how the individual looks, sounds, and talks in a dynamic persuasive communication. In fact, we show that dynamic delivery features work independently of static traits. Second, our method has high scalability and replicability. Even though the underlying algorithms are trained and cross-checked using millions of subject-rated data points, our method does not involve subject recruiting. The algorithm can be viewed as a speedy, tireless, and well-trained rater following a consistent standard; thus the method is replicable and ready to scale up computationally.

The method could potentially open new venues for research questions for economists. For example, we may study the decisions of loan officers through loan application interviews. Financial news streams can allow us to study the high-frequency impacts of financial news on financial markets. Job interview videos may allow us to study how social interactions affect labor decisions.

## I. Data and Setting

Our empirical analysis investigates venture investment decisions after startup pitches. The data set consists of two main parts—entrepreneurs’ pitch videos for accelerator applications and startups’ company- and team-level information. The two parts are manually merged using company names. The sample spans from 2010 to 2019.

---

<sup>6</sup>See, for example, [Rosenberg et al. \(1986\)](#), [Ambady and Rosenthal \(1992\)](#), [Ambady and Rosenthal \(1993\)](#), [Schubert et al. \(1998\)](#), [Todorov et al. \(2005\)](#), and [Benjamin and Shapiro \(2009\)](#). For some recent economic papers adopting the same norm, see [Berggren, Jordahl, and Poutvaara \(2010\)](#), [Brooks et al. \(2014\)](#), [Blankespoor, Hendricks, and Miller \(2017\)](#), and [Huang et al. \(2023\)](#).

### A. *Video Data and the Pitch Setting*

We use pitch videos when startups apply to five large and highly ranked accelerators in the US: Y Combinator, MassChallenge, 500 Startups, Techstars, and AngelPad. Accelerator investments are important for entrepreneurship and innovation (Hochberg, 2016; Lerner, Schoar, Sokolinski, and Wilson, 2018). As of July 2019, these accelerators have accelerated more than six thousand startups, which in total have raised over \$48 billion of total funding. Many leading entrepreneurial companies were funded by accelerators, such as Dropbox (2007), Airbnb (2009), and DoorDash (2013). Accelerator investment typically grants a standard contract with a fixed amount of investment (ranging between \$20,000 and \$150,000, fixed within accelerator-year). This allows our study to focus on one clean “yes or no” investment decision and ignore other investment parameters like amount or term sheet negotiation.

When startups apply to accelerator programs, they are required (or highly recommended) to record and submit a self-introductory pitch video of standard length as part of the application process. These videos are typically one- to three-minute long, and they present the founder(s) introducing the team and describing the business idea. These videos, rather than being submitted to the accelerators directly, are uploaded to a public platform such as YouTube and links to these videos are provided in application forms. This procedure provides researchers an opportunity to access those videos. We develop an automatic search script for two public video-sharing websites, YouTube and Vimeo. The web crawler returns a list of videos using a set of predefined search keywords, such as “startup accelerator application video”, and “accelerator application videos”, among others. Appendix A provides more details on this process.

**[Insert Table I Here.]**

This process yields 1,139 videos used in our analysis. In Table I Panel A we report basic information at the level of video pitches. The median length of a pitch video is 68 seconds, and the mean is 83 seconds. These videos are not viewed often, and most do not attract any likes or dislikes. This is consistent with the fact that these videos are generic pitch videos for application purposes rather than for any marketing campaign or product promotion. Regarding team composition, we find that 46 percent of the startups have only one member pitching, and 54 percent have multiple members. The average number of members per video is 1.74. Forty-nine percent of the teams have only male founders, 27 percent have only female founders, and 24 percent have mixed genders.

We want to note that the videos in our analysis are an incomplete sample of all pitch videos ever made by accelerator applicants. Many startups may have chosen to unlist or privatize their videos to make them unavailable to us, as researchers, to search and view. We formally discuss the sample selection issue and its implications for our analysis in Section III.A.1, and we show that this sample selection does not affect our findings or interpretations. Moreover, we want to acknowledge that our sample focuses on a group of early-stage venture investors whose investment decision-making may not be completely representative of the venture investment community.

In our setting, investors did not view these pitches in person—these pitches are recorded and uploaded for investors to review during decision-making. As a result, when interpreting our findings, one needs to be mindful of how the same features (e.g., smile, passionate voice, word choices) that affect decision-making when watching the video can translate to in-person interactions. Reassuringly, [Dana, Dawes, and Peterson \(2013\)](#) show that experiencing in-person interviews and watching video interviews lead to similar biases. Using videos also facilitates the connection to empirical studies in other fields on persuasive materials, like advertisements, media materials, etc.

### *B. Startup Information and Team Background*

We also collect startup information on both the companies and the founders. In venture investment, investors value human assets like education and work experience ([Bernstein, Korteweg, and Laws, 2017](#); [Howell, 2019](#)). They may also be, often mistakenly, influenced by discriminative factors, most noticeably gender ([Gompers and Wang, 2017](#); [Ewens and Townsend, 2020](#); [Gornall and Strebulaev, 2022](#); [Hebert, 2020](#)). We incorporate those investment determinants in our analysis.

Startup information is collected from two widely used entrepreneurship databases, Crunchbase and PitchBook. We start by searching for companies in these two databases according to the names identified in application videos using video titles, subtitles, and uploader ID. For startups with duplicate names or name changes, we identify companies by names of founders, business descriptions, and time of founding. Startup-level variables include the year of founding, location, operating status, total funding round and amount, number of investors, and number of employees.

Among the 1,139 applications, 462 are listed in Crunchbase, 217 are in Pitchbook, and 208 appear in both databases. Thus, 471 unique ventures are covered by at least one of the VC databases. Panel B of Table I reports the summary statistics for these startups. Nine percent of the startups in our sample received funding from accelerators. As of July 2023, the average firm age is 6.2, with

a median age of 6. Regarding long-term performance measures, and after setting these measures to zero for failed startups, we find that the mean number of employees is 8. Ten percent of the startups raised VC funding, and three percent reached a milestone exit event such as an IPO or acquisition. The average annual frequency of their website updates over the three-year period after the accelerator applications, as tracked using the Wayback Machine, is 0.8 times. We also report these summary statistics for a subset of 270 startups that received seed investment from an early-stage investor. Their average performance measures are better than that of the full sample.

We also standardize companies’ industry classifications using the Global Industry Classification Standard (GICS). We categorize all companies into one of 24 GICS industry groups, which then form 11 GICS sectors, using the industry information in Crunchbase and PitchBook along with the video scripts.

Beyond company-level information, we also collect information on the founding teams. We compile a list of founders for each startup company using Crunchbase, PitchBook, and video content. For each startup team member, we use LinkedIn to extract the five most recent educational experiences and the ten most recent work experiences. This information is used to construct variables that indicate each presenting team member’s education (university, degree), job seniority, entrepreneurship experience, etc.<sup>7</sup> Among 1,139 startups in our sample, we are able to find the presenters on LinkedIn for 693 of them. For those that we are unable to find LinkedIn profiles, we code corresponding variables as missing and include the video in our analysis. As shown in Table IA.3, our main results are robust on this sample of 693 companies. We report the summary statistics of team composition and background in Panel C of Table I. Regarding the startup founder teams’ backgrounds, we find that 49% of the founders have startup experience at the time of application, conditional on finding their LinkedIn profiles. Thirty-two percent and 5% of them hold Master’s degrees and PhD degrees, respectively, conditional on finding their LinkedIn profiles.

## II. Method: Processing Video Data with ML Algorithms

Our video processing method proceeds in three steps. First, we decompose the information embedded in the videos into three-V dimensions—visual, vocal, and verbal. Second, for each dimension, we adopt ML algorithms to create visual, vocal, and verbal features from the raw data.

---

<sup>7</sup> When constructing team-level controls, we only focus on team members who are present in the pitch videos. Our empirical results are not sensitive to this decision. See Internet Appendix A.2 for more details on data collection.

Third, we aggregate these measures within and across these dimensions to characterize features representing each pitch video. Below we first give an overview of the key properties of video data relevant to economic research, and then we describe the three-step method in depth. Appendix B provides additional technical details.

It is worth noting that the method is intentionally built upon off-the-shelf algorithms and applications developed in the past decade by computer scientists and reputable cloud computing providers instead of being built and trained by us. The key motivations behind this decision are to allow easy replication and easy adaptation, to maximize transparency, and also to lower the hurdle for other researchers to implement the method.

#### A. *Video as Data: Key Properties*

Videos are a pervasive form of data. More than 80 percent of the world’s internet traffic consists of transmissions of video, and more than 60 percent of the total digital data stored worldwide are video. However, videos are underexploited in economics research largely due to the complexity and computational burden. We begin by discussing some basic properties of videos and illustrate how these properties relate to our video processing and measure construction process.

First, video data are information intensive. To better understand the richness of video data, one can make a size comparison between video files and other data files. The csv-format startup panel in this study is around 1 MB in size. It includes 150 company-level characteristics for 1,139 startups. In contrast, a one-minute high-resolution pitch video in mp4 format can be as large as 200 MB in our sample. To put this into perspective, one second of a high-definition video, in terms of size, is equivalent to over 2,000 pages of text (Manyika et al., 2011; Gandomi and Haider, 2015).

Second, video data are unstructured and high-dimensional, making them more complicated to process relative to other data formats such as panel data or even unstructured textual data. Consider a one-minute video with a resolution of  $1280 \times 720$  (720p) and two 48 kHz audio channels. In this case, there are  $1280 \times 720 = 921,600$  pixels in each image frame. If we use a sampling rate of ten frames per second, the video can be represented as a series of 600 images. These 600 images need 552 million pixels (dimensions) to be represented. Further, to represent the acoustic information, we need  $48,000 \times 2$  dimensions per second and thus, around 5 million dimensions for one minute. In total, to represent such a video, we need around 560 million dimensions.

Third, video data have a low signal-to-noise ratio (SNR) and low economic interpretability.

Regarding the SNR, most videos have a large amount of background noises that are irrelevant to the primary economic question. In our setting, these noises include background noise, and furniture in the video, among other things. Moreover, the information units of video data (e.g., the pixels and the sound waves) are not directly interpretable as accounting variables or textual words. Thus, when processing video data, we need to impose structures and algorithms to guide the extraction of information that is useful and meaningful for economic research.

### *B. Step 1: Information Structure and Representation of Video Data*

The first step in video data processing is to decompose videos into the three-V structure and to represent them in a data format. Three-V means visual, vocal, and verbal; this structure is intuitive and widely accepted in the social psychology and communication literature (Mehrabian, 1972; Strahan and Zytowski, 1976; Krauss et al., 1981; Knapp et al., 2013).

We first extract these three dimensions from the two streams of digital information in raw video records—the image stream and the sound stream. For the visual component, we represent the video using images sampled at ten frames per second and employ face-detection ML algorithms to identify human faces in each video frame.<sup>8</sup> The unit for visual analysis is thus at the level of each video frame. For the vocal component, we extract the audio files from the video. For the verbal component, a speech-to-text ML algorithm is used to extract speech from the sound.

We now consider each of these three dimensions and discuss their data representations one by one.<sup>9</sup> For visual, human faces exist in the format of digital images, which can be numerically represented as two-dimensional matrices. Moreover, since we have a stream of such digital images, the information in the facial dimension is coded as a time series of two-dimensional matrices. The vocal signal, essentially sound waves, can also be digitized as a time series of amplitudes given a specific sampling rate (the number of times per second that the amplitude of the signal is observed) and resolution (the number of discrete levels to approximate the continuous signal). If the audio is multi-channel, it can be represented as a matrix with channels as columns/rows. Finally, we transform human speech into a term-frequency matrix, which tracks the use of words in the verbal

---

<sup>8</sup>0.1 seconds is a very short time interval for our study and for the goal of capturing facial expressions. Even fast facial movements, like blinking, take on average around 0.25 seconds, and will be captured by our algorithm at the 1/10th-second intervals.

<sup>9</sup>This step builds on and extends the `pliers` package available at <http://github.com/tyarkoni/pliers> and McNamara, De La Vega, and Yarkoni (2017).

content.

### *C. Step 2: Constructing Measures with Machine Learning*

In the second step, we construct economically interpretable measures from the represented video data using ML algorithms. One way to think about our ML algorithms is as a robust and objective super robot that can rate and record the video data at a high frequency along the three-V dimensions. The ML algorithms are trained using millions of observations rated by subjects and are now automatic and easy to scale up. To do so, we leverage many recent groundbreaking advances in computer vision, speech recognition, and textual analysis.

#### *C.1. Visual (Image)*

We identify human faces embedded in each video image using face detection algorithms. For replicability and transparency, we directly adopt the established implementation of Face++.<sup>10</sup> The Face++ platform provides APIs through which we feed our raw images into the cloud computing system and receive a host of face-related measures constructed by Face++’s ML algorithms. We also check the robustness using an alternative computation platform, Microsoft Azure Cognitive Services.<sup>11</sup>

The process is as follows. First, the algorithm detects locations of facial landmarks (e.g., nose, eyes) from the raw images using face detection technology. These coordinates allow us to detect facial movements, such as smiles, eye blinks, or the lowering of eyebrows. They also detect mouth movements that can help identify the speakers continuously. Second, this information enters emotion recognition algorithms that categorize facial emotions into one of the following six dimensions—happiness, sadness, anger, fear, disgust, and neutral. In the empirical analysis, we combine the measures of sadness, anger, fear, and disgust into a composite facial negative emotion measure. We exclude neutral facial emotion in our analysis because the sum of facial positiveness, negativeness, and neutrality is one, which induces collinearity. Third, we obtain a face beauty measure and some demographic characteristics of individuals, including gender and predicted age.

These variables are different from the static features of beauty and impressions (attractiveness,

---

<sup>10</sup>Face++ can be accessed at: <https://www.faceplusplus.com/emotion-recognition>.

<sup>11</sup>Microsoft Azure Cognitive Services can be accessed at: <https://azure.microsoft.com/en-us/services/cognitive-services/face>.



competence, etc.) that are used in prior economic studies. Intuitively, we take out the static features (e.g., people can be more or less attractive) that can be viewed as baseline features. We then track the movement of the face to detect facial emotions (e.g., passionate—which can be a feature of speakers with both more and less attractive appearances).<sup>12</sup>

*Example.* In Figure IA.3 in the Internet Appendix, we present sample frames of high-positivity and low-positivity facial expressions. As discussed earlier, the calculations of positivity, negativity, and other facial features are done at a frequency of every 1/10 of a second.

### C.2. Verbal (Text)

Next, for verbal, we extract human speech from audio data using the speech-to-text conversion API provided by Google Cloud.<sup>13</sup> This ML-based API converts audio into a text transcription. These transcripts include a list of words, time stamps (onsets, offsets, and durations) of these words, and punctuation. We then merge the speech corpus with two dictionaries. The first dictionary is the Loughran-McDonald Master Dictionary (LM hereafter), which is commonly used in financial text analysis and provides text categories such as negative and positive, among others (Loughran and McDonald (2011)). The second dictionary is developed by social psychologists using Wordnet and word embeddings (Nicolas, Bai, and Fiske (2019)). This dictionary (NBF hereafter) includes word categories along the dimensions of social psychological traits such as ability and warmth, which helps us to measure verbal characteristics from the angle of social perception (Fiske et al. (2007)).

*Example.* Given that researchers in economics and finance are already familiar with the negativity and positivity categorizations in textual analysis, we here provide an example of the “ability” and “warmth” dimensions in verbal expressions. In Figure IA.4 in the Internet Appendix, we show a high-ability script in Panel A and a high-warmth script in Panel B. High-ability pitches focus on the ability of the entrepreneur and the operational efficiency of the business idea. In contrast, high-warmth pitches focus on communion, pleasing personalities, and a description of a bright future working with VCs and customers.

---

<sup>12</sup>This image processing technology is related to some recent work that explores images as data (Joo and Steinert-Threlkeld, 2018; Peng, Teoh, Wang, and Yan, 2021), particularly in communication and political science (Peng, 2018; Boxell, 2024). These papers typically focus on static images. Our images are continuous and extracted from videos, and the processing is conducted for the purpose of capturing dynamic interaction features from the video.

<sup>13</sup>Google Cloud Speech-to-Text API can be accessed at: <https://cloud.google.com/speech-to-text>.

### C.3. Vocal (Voice)

Finally, we analyze information embedded in the vocal channel. In this part, we regard voice as digital signals and focus on the physical information not captured by textual transcription of human speech. Different from images that can provide rich information independently, audio data, essentially sound waves, code information in the audio’s dynamics and auto-dependence structure. In other words, if we split the audio into fixed high-frequency segments like a series of images, we may lose the information embedded in the auto-dependence structure. We address this technical problem by focusing on word or sentence units by leveraging the outputs of speech-to-text algorithms. Specifically, we split each audio stream into segments by words and sentences. These units naturally reflect the auto-dependent information structure in the video and are also good approximations of human cognitive processes.

We employ `pyAudioAnalysis` (Giannakopoulos (2015)), which is a Python package for audio feature extraction, classification, segmentation, and application. We extract 34 low-level features in total. These features include but are not limited to spectrogram, chromagram, and energy. These features capture the physical characteristics of the vocal channel.<sup>14</sup>

We then construct high-level cognitive measures by feeding these low-level audio features into ML algorithms performing vocal emotion analysis. We adopt two conceptual frameworks for vocal emotion. The first framework models vocal emotion along two dimensions—arousal and valence (Bestelmeyer et al. (2017)). Valence measures the positive affectivity of the vocal feature, while arousal captures the strength of such an attitude (exciting versus calm)—in some sense, it captures the concept of “being passionate” often referenced in entrepreneurship studies (Gompers et al. (2020)). We use `pyAudioAnalysis`’s established implementation and pre-trained ML models to obtain vocal arousal and valence. Another framework models vocal emotion along three emotional dimensions—happy, sad, and neutral. To implement this conceptual model, we use `speechemotionrecognition`, which includes deep-learning-based speech emotion recognition algorithms. We adopt the pre-trained ML models in this package and feed our audio segments into these models.<sup>15</sup>

*Example.* In Figure IA.5 in the Internet Appendix, we provide the waveform amplitude plot for

---

<sup>14</sup>For a complete list of these audio features, please check <https://github.com/tyiannak/pyAudioAnalysis/wiki/3.-Feature-Extraction>.

<sup>15</sup>The deep-learning models in this package are trained on the Berlin Database of Emotional Speech (Burkhardt et al. (2005)). Vocal-Neutral is dropped from our analysis due to collinearity.

two pitches, one with high arousal (i.e., high excitement) and one with low arousal. The general patterns of the sound waves differ significantly and can be coded by our ML algorithm. Even though it is more difficult to visualize other features, the logic applies—analyzing the sound waves allows the categorization of vocal features.

#### D. Step 3: Measurement Aggregation

After constructing measures for each channel separately, we merge these measures and aggregate them to the video level. An appendix table of measures constructed in this paper along with their definitions and construction algorithms/procedures can be found on page 41.

##### D.1. Simple Aggregation

We first aggregate all measures at the video level by taking the mean of measures across everyone in the pitch and across the whole video. We create *Visual-Positive*, *Visual-Negative*, *Vocal-Positive*, *Vocal-Negative*, *Vocal-Valence* and *Vocal-Arousal* to capture facial and vocal emotions. For each of these variables, we compute the average proportion of time in the pitch that a team member shows certain facial or vocal emotion. We create *Visual-Positive* as the score of visual happiness and create *Visual-Negative* as the combined score of visual anger, sadness, fear, and disgust. This allows us to mitigate potential measurement errors introduced by the ML algorithm in classifying subtly different negative visual emotions. For vocal measures, we label vocal happiness as *Vocal-Positive* and vocal sadness as *Vocal-Negative*. For verbal content, we focus on *Verbal-Positive* and *Verbal-Negative* in the LM financial dictionary and *Verbal-Warmth* and *Verbal-Ability* in the NBF social psychology dictionary. *Verbal-Positive* and *Verbal-Negative* are calculated as the word counts (in each category) scaled by the total number of words in each pitch. *Verbal-Warmth* and *Verbal-Ability* are calculated as the directed word counts (+1, if in the positive direction of the category; −1, if in the negative direction; 0, if unrelated to the category), which are also scaled by the total number of words. Figure 1 illustrates the data structure and shows how we transform unstructured video data into our final structured panel data set.

**[Insert Figure 1 Here.]**

Table II presents the summary statistics for pitch videos. We show in Panel A the mean, median, and 25th and 75th percentiles. There are substantial cross-sectional variations in startup pitches

along the three-V features. Take *Visual-Positive* as an example. The mean, 0.17, can be roughly seen as indicating that in a pitch video, on average, the speakers show clear happy visual features about 17 percent of the time. But this number varies quite dramatically. At the 25th percentile, one speaker could show happy facial features only 5 percent of the time, while the 75th percentile most positive team could score 25 percent in this measure. Some features, especially those capturing negativity, have low means. For instance, the vocal and verbal negativity measures both score 1 percent at the mean—this is not surprising given that entrepreneurs would likely try to hide negativity during pitches. This does not mean, however, that those negativity measures are less meaningful—on the contrary, as will be shown below, the negativity features are important in our analytical results.

**[Insert Table II Here.]**

In Table II Panel B we show the correlation between metrics from different channels. We find that within a given dimension (the same V), features are highly correlated—for example, videos showing more positive attitudes naturally will show fewer negative attitudes.<sup>16</sup> Meanwhile, across dimensions, similar metrics correlate in very interesting ways. Positivity in vocal features is positively correlated with positive visual expressions, confirming the validity of our metrics that are actually generated using completely different information and algorithms. But the correlation is mild, suggesting that vocal and visual expressions are correlated yet separate signals. Verbal information is uncorrelated with vocal and visual features. This could be because textual scripts can be more easily prepared and recited, and thus they can be disjointed from the vocal and visual delivery.

#### *D.2. Generating a “Pitch Factor”*

Beyond the set of detailed features, one may wonder—can we create one variable to summarize “how well” entrepreneurs deliver their pitches? We achieve this goal using factor analysis to extract the most important common component from the variance–covariance matrix of the features in pitch videos. This process allows us to eliminate the redundant features in the complex pitch structure.

Operationally, we estimate the factor analysis using the principal component method similar to Tetlock (2007). The method chooses the vector in the pitch feature space with the greatest

---

<sup>16</sup> The outliers are correlations among verbal measures, which are low. This is potentially due to the fact that the widely adopted dictionaries categorizing sentiment cover only two percent of all the words used in business documents (Loughran and McDonald (2011)), and 98% are considered as neutral words.

variance, where each feature is given equal weight in the total variance calculation. We explore other factor analysis estimation methods, such as principal factor analysis and maximum-likelihood factor analysis. The qualitative empirical conclusions are not sensitive to the method chosen, and the quantitative conclusions change only minimally. Effectively, principal components factor analysis performs a singular value decomposition of the correlation matrix of pitch features. The single factor selected in this study is the eigenvector in the decomposition with the highest eigenvalue—we label it as the “Pitch Factor.”

The Pitch Factor not only summarizes the pitch delivery but bears a clear and intuitive interpretation. The last two columns of Table II Panel A report each variable’s loading on the Pitch Factor and its “uniqueness”. The loadings are positive for all measures that have positive and affirmative economic meanings. For example, the factor loads positively on *Visual-Positive* (+0.08), *Vocal-Positive* (+0.39), *Vocal-Arousal* (+0.91), *Vocal-Valence* (+0.88), *Verbal-Warmth* (+0.06), and *Verbal-Ability* (+0.06). Meanwhile, the factor loadings are negative for all measures in the negative direction, such as *Visual-Negative* (−0.14), *Vocal-Negative* (−0.30), and *Verbal-Negative* (−0.14). Together, the Pitch Factor can be viewed as a composite index that integrates the information from three-V channels and represents the overall level of positivity, passion, and warmth reflected in the pitch video. The uniqueness is the percentage of variable variance that cannot be explained by the factor. The uniqueness is low on average, so the factor is well-behaved and powerful.

### D.3. Cross-Validation with Human Raters

As a way of cross-validation, we compare our Pitch Factor with the ratings solicited from respondents at Amazon Mechanical Turk (the MTurkers). Appendix D provides a detailed description of the survey designs, and here we overview the exercise only briefly in order to avoid distractions from the main method.

We use two surveys to document the high correlation between our Pitch Factor and the ratings provided by MTurkers. In the first design, we directly elicit ratings from participants to compare with Pitch Factor. The MTurk survey follows previous research to validate ML-based measures using human raters (Peng et al., 2021).<sup>17</sup> We recruited 115 respondents to watch pitch videos and provide ratings on the overall level of positivity in the pitches on a scale of 1-9, where positivity

<sup>17</sup>MTurk is increasingly used for other purposes in economic research (DellaVigna and Pope, 2018; Lian, Ma, and Wang, 2019).

is defined as enthusiasm and passion for the respondents. We show in Figure [IA.6](#) and Table [IA.8](#) that a strong correlation exists between the Pitch Factor generated from our ML-based method and the human-rated positivity score, even after controlling for rater FE. In our second design, we ask MTurker respondents to compare pitch positivity for pairs of randomly-drawn videos. For each of these random pairs, we evaluate the consistency between our ML-based ranking and the human ranking. In other words, does the algorithm pick the same winners as the raters? We find that the same winner is picked with nearly 89.5% consistency.

### III. Empirical Analysis

This section presents our main empirical analysis and proceeds in three steps. We first show the role of the Pitch Factor in explaining venture investment decisions. We also discuss the robustness of the results accounting for the textual content of pitches, sample selection concerns, and omitted variables. Next, we examine whether the “positive pitch driving investment” pattern helps investors make better decisions. We study the post-investment performance of the startups and show that startups with more positive pitches, conditional on obtaining funding, underperform. At last, we further explore the potential biases involved in this process by exploring an important source of heterogeneity—the gender of founders.

#### A. Baseline Result: Positivity Pitch Features and Venture Investment

Our first analysis examines whether startup  $i$ ’s Pitch Factor relates to its likelihood of obtaining funding when applying to accelerator  $j$  during application year  $t$ . This is a cross-sectional data set with 1,139 pitches since each investment evaluation happens only once in the sample. The analysis is performed using the following specification:

$$I(Invested)_{ijt} = \alpha + \beta \cdot \text{Pitch Factor}_i + \gamma \cdot \text{Controls}_i + \delta_j + \varepsilon_{ijt}. \quad (1)$$

The key outcome variable is  $I(Invested)$ , which equals 1 if the startup was chosen by the accelerator and 0 otherwise. On the right-hand side, the Pitch Factor is standardized into a zero-mean variable with a standard deviation of one.

The model controls for pitch content, for example, if the idea sounds novel or if the pitch covers

important dimensions of the business model. We measure each pitch’s content using the textual script and construct three sets of variables.<sup>18</sup> First, we measure the novelty of ideas in video pitches based on their textual similarities with other existing businesses, where the similarities are calculated using a BERT (Bidirectional Encoder Representations from Transformers) model. To implement this, we compare pitch scripts to business descriptions of startups from PitchBook and public firms from 10-K filings. Following [Kelly et al. \(2021\)](#), who use patent documents to measure the novelty of patents, we categorize an idea as novel if it differs from existing business ideas in the economy but is similar to successful companies in the near future.

Second, we take a dictionary-based approach and compile a word list covering topics that are often considered as important for startup financing (see Table [IA.4](#) in the Internet Appendix for the word list). These categories include cash flow, employment, readiness, technology, data and AI, competition, and concrete numbers. For each video pitch, we construct an array of dummy variables indicating whether each category is mentioned in the pitch.<sup>19</sup>

Third, as an extension to our dictionary-based approach, we also use word categories in Linguistic Inquiry and Word Count (LIWC). LIWC is a textual analysis software widely used in computational linguistics and finance ([Kuhnen and Niessen, 2012](#); [Gow et al., 2016](#); [Braggion et al., 2017](#); [Kogan et al., 2021](#); [Gómez-Cram and Grotteria, 2022](#); [Farrell et al., 2022](#)). We focus on two sets of categories: communication styles (e.g., concrete or informal words) and time orientations (e.g., past, present, or future focus).

We control for accelerator fixed effects to account for the possibility that certain accelerators might attract specific types of startup founders or have different investment criteria or preferences that could correlate with pitch features. Standard errors are clustered at the accelerator-year level. This accounts for the fact that an investment decision for one startup is automatically correlated with the accelerator’s decisions about other startups applying in the same year, given the investment quota constraint that accelerators face.<sup>20</sup>

**[Insert Table [III](#) Here.]**

---

<sup>18</sup>Appendix [B.3](#) provides additional details on these measurement constructions.

<sup>19</sup>We also construct measures for the frequency of words in each category appearing in a pitch. The results are not sensitive to this decision.

<sup>20</sup>Appendix Table [IA.9](#) shows that the results are qualitatively identical and quantitatively similar when using visual measures constructed using Microsoft Azure instead of the Face++ API. Table [IA.10](#) shows that the results are robust to alternative fixed effects combinations such as industry fixed effects and to estimation methods such as OLS. The results are also robust to controlling for video resolution, image brightness, color intensity, etc.



Table III presents the results of Eq. (1) estimated using a logit model, reporting marginal effects estimated at the sample mean. Column (1) includes only the Pitch Factor, which captures the pitch’s overall level of positivity, passion, and warmth. The Pitch Factor positively and strongly correlates with the probability of the startup receiving funding from the accelerator. The 0.030 coefficient means that a one standard-deviation increase of the factor is associated with a change of three percentage points in funding probability, which is equivalent to a 35.2 percent increase from the baseline funding rate of 8.52 percent.

This large economic magnitude shows the importance of non-verbal features in persuading investors and is in line with previous research. For example, Kaplan and Sorensen (2021) find that interpersonal skills increase the probability of an executive being hired by about 20% (based on numbers from Table VII of the paper). Gorodnichenko et al. (2023) demonstrate that making one’s tone of voice in FOMC meetings sound more positive by one standard deviation could raise S&P 500 returns by approximately 75 basis points, which is as powerful as a similar change in forward guidance, the explicit communication about future monetary policy by the Fed.

We also interpret this economic magnitude through the persuasion rate (DellaVigna and Kaplan (2007)). Consider the following thought experiment: judges view a pitch with either a high pitch factor (top twenty percent) or a low pitch factor (bottom twenty percent). The former is labeled as treated, by the passionate pitch, and the latter as untreated. In our sample, the probabilities of receiving funding in the treated and control group are 12.33 percent and 6.14 percent, respectively. The persuasion rate is calculated as<sup>21</sup>

$$f = \frac{I_{HighPitchFactor} - I_{LowPitchFactor}}{1 - 0} \cdot \frac{1}{1 - I_{Baseline}} = \frac{12.33\% - 6.14\%}{1 - 0} \cdot \frac{1}{1 - 6.14\%} = 6.60\%. \quad (2)$$

Essentially, the approach scales the change of investment probability after being exposed to the treatment (the first fraction) by the effective room for persuasion after excluding the baseline investment rate, which is unobserved and conventionally approximated using control group behavior (i.e.,  $I_{Baseline} = I_{LowPitchFactor}$ ).

Columns (2)–(5) control for informational content measures, at first one category at a time and then all in the same regression. Controlling for informational content measures barely moves the magnitude and statistical significance of the Pitch Factor’s coefficient, indicating that the Pitch

---

<sup>21</sup>For a more general discussion on persuasion rates, see DellaVigna and Gentzkow (2010).

Factor remains an important and independent feature in determining funding outcomes. In addition, some of the content controls significantly correlate with the probability of obtaining funding. More noticeably, in column (2) the business novelty measures (Idea Novelty (PB) and Idea Novelty (10K)) strongly predict funding probability. As shown in column (3), concrete discussions on “cash flow” are also related to higher funding probability. Column (4) shows that the concreteness of the pitch, and the discussion on past success rather than future outlooks, are associated with higher funding probabilities.

**[Insert Table IV Here.]**

Table IV presents the results of Eq. (1) estimated for both the Pitch Factor and individual visual, vocal, and verbal pitch features. All the features are standardized into a zero-mean variable with a standard deviation of one, making the economic magnitudes easier to interpret and compare across variables. For each feature, we show the marginal effect calculated at the sample mean, the standard error, and the pseudo  $R^2$  in each row. We present both models without startup/team controls (left panel) and with comprehensive controls (right panel).

The visual, vocal, and verbal measures present a consistent message as the Pitch Factor. A one-standard-deviation increase in happiness reflected in the visual dimension is associated with a 1.5 percentage point increase in investment likelihood, or a 17.6 percent increase from the unconditional funding rate. Startup teams showing more negative facial expressions, denoted as Visual-Negative, are less likely to receive funding. The absolute economic magnitude is 93 percent larger than that of Visual-Positive (0.029 versus 0.015), suggesting that the negative component is even more relevant in driving investment decisions. This resonates with findings that the negative spectrum is more relevant in research on the “beauty premium” (Hamermesh and Biddle (1994)) and textual “market sentiment” (Tetlock, 2007; Loughran and McDonald, 2011).

Note again that the visual delivery measures are captured dynamically and are independent of static facial traits such as beauty (and similar attractiveness, competence, etc.). To get a sense of this independence and a benchmark for understanding the above economic magnitude, we lean on the well-established concept of the “beauty premium,” which is shown to be important in the labor market and other economic decisions (Hamermesh and Biddle, 1994; Mobius and Rosenblat, 2006; Graham et al., 2016). We confirm that more attractive entrepreneurs are more likely to receive investment. The economic magnitude of the beauty effect is roughly the same as it is for

Visual-Positive and is smaller than for Visual-Negative.

For the vocal dimension, we rely on two vocal emotion categorizations. In the first positive-negative categorization, the pattern is quite similar to the findings on visual features. More positive (negative) tones in pitches are associated with a higher (lower) probability of receiving accelerator financing. In the second categorization, the audio channel is projected to a two-dimensional space of arousal and valence. We find that high-valence and high-arousal pitches are more likely to attract investment: a one-standard deviation change in either measure is associated with an increase of 2.1 percentage points (24.6 percent from baseline) in the probability of receiving funding.

Regarding the verbal dimension, using positive versus negative words in pitches matters for the investment decision. Consistent with previous research using this categorization in economics, negative words are more relevant for economic outcomes ([Loughran and McDonald \(2011\)](#)). Verbal-Warmth and Verbal-Ability dimensions are based on social psychology dictionaries and capture how each word influences the listener’s perceptions. Warmer pitches (i.e., friendlier and happier) attract more investment, while teams that talk more about their ability and competitiveness more often drive investors away. This is somewhat surprising given that entrepreneurial investment is a professional decision involving identifying more capable entrepreneurs.

Overall, the baseline result suggests that non-content delivery features in visual, vocal, and verbal channels all matter for financial investment decision-making in persuasive communication. To test the robustness of our method and the results, we replicate our analysis using a completely different sample compiled from an incubator program run by a reputable US university (“university sample”). We find qualitatively and quantitatively similar results (see Internet Appendix [C](#)).

Below, we further discuss this main specification.

### *A.1. Sample Selection of Videos*

Pitch videos that can be accessed and collected from our internet search are a subsample of all such videos in accelerator applications. This is because some video files are made private, unlisted, or removed from the hosting websites after the application. The empirical regularity that governs the video selection process would naturally affect the validity and accuracy of the findings thus far. For instance, if pitch videos available to researchers are selected based on having more positive features and simultaneously are more likely to be kept available by invested startups, our main results could be driven by the selection mechanism.

We address this issue by directly tracking sample selection. Our goal is to explore the primary concern—whether the selection is related to the pitch features and to whether the startup gets invested in. All the videos used in this paper were available (searchable and viewable) in July 2019. We focus on all 527 videos uploaded within the 18-month window prior to July 2019, that is, the 2018 and 2019 cohorts, since attrition is more active immediately after the uploading and program selection. By the end of March 2020, 126 videos, or 23.9 percent, were selected out (unlisted, privatized, or removed) from the hosting platforms.

**[Insert Table V Here.]**

Table V shows the determinants of the selection. The selection, or equivalently the selection out, is unrelated to the Pitch Factor, the investment decision of the accelerator, or broadly any future VC financing. All coefficients are statistically weak and economically minimal across different specifications of the selection model. For instance, the 0.006 coefficient in column (1) means that a one standard deviation change of the Pitch Factor shifts the selection probability by only 0.6 percent points, or 2.5 percent from the baseline, with a wide standard error. Thus, this type of sample selection can be considered quasi-random for our study.<sup>22</sup>

Another layer of the sample selection problem arises from the initial decision to submit and publicize a video pitch. Since we do not observe the administrative data from the accelerators, we cannot test this directly in our context. However, we tackle this selection dimension using the university sample previously mentioned and documented in Appendix C. In this sample, we observe administrative-level data of submitted applications and their video links (instead of us scraping when forming our main sample). We show in Table IA.6 that our main results remain robust both statistically and economically in this sample, suggesting that the potential sample selection issues associated with video availability are unlikely to be a first-order concern in our study. We also perform a selection model using the university sample and find little evidence (as shown in Table IA.7) that the availability of a video in the initial stage is related to measures of pitch features or investment decisions. Even though the analysis is limited by the size of the sample and statistical power, it assures us that this layer of the selection issue may not be a major concern.

---

<sup>22</sup>One caveat is that the selection analysis conditions on a video once being publicly available at some time. This leaves the possibility that some videos were never made public. However, the selection trade-off at the initial uploading should be fairly comparable to the selection model estimated in Table V.

## A.2. *The Value of the Video-Based Method*

Leveraging ML techniques, our video processing method has two distinct features from the prior literature that are particularly valuable in studying persuasion delivery. First, we use the complete pitch video when constructing the measure, differing from the literature that exploits thin slices of data (like the first few seconds) to capture perceptions. Second, we jointly use information from the three-V dimensions, raising curiosity regarding the value of our method over those earlier works that examined some dimensions individually.

**[Insert Table VI Here.]**

How valuable are those two features? The general design of the tests is to horse-race the Pitch Factor with measures constructed using thin-sliced video clips and with those using single-V dimensions. We report the analysis in Table VI. In Panel A, we compare our Pitch Factor with factors constructed using thin slices of data. First, we construct two new Pitch Factor measures, one using only a clip of when the first word is spoken (roughly one second), and the second using a random second of video clip in the pitch. In columns (1) and (3), we find that thin-sliced measures do provide useful information in explaining the investment choice. However, in columns (2) and (4), once we incorporate the full-video Pitch Factor to horse race with the thin-sliced versions, the full-video Pitch Factor dwarfs the thin-sliced versions. In column (5), all three variations of the Pitch Factor are incorporated in the same analysis—again, the Pitch Factor using the full video dominates.

In column (6) of Panel A, we present another analysis to demonstrate the marginal informational gain from using the full video in constructing the Pitch Factor. We implement a Shapley-Owen decomposition (Shorrocks (2013)) to show that Pitch Factor contributes the largest fraction of  $R^2$  when explaining investment decisions—significantly larger than the first-slice version and the random-slice version. Using the same specification as in column (5), we decompose the total (pseudo)  $R^2$ , and the results are shown in column (6).<sup>23</sup> We find that among all the  $R^2$  that can be explained by Pitch Factors, the full-video Pitch Factor has the largest  $R^2$  contribution (nearly 67%), which is four times larger than that attributable to the first-slice and random-slice Pitch Factor. In

---

<sup>23</sup>In a nutshell, this method calculates the contribution of each covariate to the (pseudo)  $R^2$  of our investment decision logit regression. It accounts for complementarities among covariates by calculating the Shapley value by averaging over the marginal contribution to  $R^2$  for every possible covariate combination.

unreported results, we also find that the  $R^2$  contribution of the full-video Pitch Factor is the largest among all the included variables, higher than that of commonly viewed important factors such as whether the startup members have prior entrepreneurial experience.

In Panel B, we investigate the value of using the Pitch Factor over the information from individual V-dimensions. In our simplistic approach, we apply the factor analysis approach to each of the three-V dimensions and construct vocal-, visual-, and verbal- factors. We find that even though those individual dimensions matter in the decision-making, they are later dominated by the Pitch Factor which aggregates all various channels. In other words, jointly using the three-V channels indeed provides a more comprehensive reflection of the pitch delivery.

The factor analysis method used in Pitch Factor construction, though simple by design, accounts not only for the three-V channels but also, importantly, the covariance matrix of information in those channels. A natural implication is that our finding is less prone to be confounded by the omitted variable bias than those using a single dimension alone. For example, if positive facial expressions are positively correlated with other features of delivery (e.g., passionate voices), it is hard to correctly estimate the effect of facial expressions when only analyzing the facial information and without accounting for other channels.

#### *B. Is Omitted Startup and Founder Quality Driving The Results?*

The ability to deliver pitches is not randomly allocated—it may be affected by education and experience, among other factors. Is the documented impact of persuasion delivery just due to the omitted quality proxies for the startup that are not explicitly controlled for in the baseline analysis?

Building on the statistical discrimination literature, we can add control variables that are good proxies of startup and team quality and track the stability of coefficients associated with pitch features. [Oster \(2019\)](#) suggests a test for omitted variable bias that uses the information contained in the change in coefficient and the change in  $R^2$  when moving from uncontrolled to controlled regression. The basic intuition is that if the coefficients are stable as we add (good but imperfect) quality controls, then the estimated effect is probably not due to an omitted quality variable and should be interpreted as arising through other independent channels. Formally, [Altonji, Elder, and Taber \(2005\)](#) and [Oster \(2019\)](#) show that if selection on the observed controls is proportional to the selection on the unobserved controls, then we can compute an identified set and test whether the identified set for the treatment effect includes zero.

We repeat the analysis in Table IV, adding control variables for founders' education (whether they have a master's degree or PhD degree and whether they attended elite universities, as defined by U.S. News & World Report's Top 10) and founders' entrepreneurship experience, prior work experience, and gender. These variables cover a large set of the information that the investor sees in addition to the pitch, and these variables have been shown to correlate with entrepreneurship quality, success, and the probability of obtaining venture financing (Bernstein et al., 2017; Ewens and Townsend, 2020).

The results are reported in the right panel of Table IV. All that we learned from the no-additional-control regression remains statistically robust. But perhaps more importantly, the coefficients remain very stable after introducing controls that are likely correlated with team quality. The formal test incorporates the change in  $R^2$  induced by adding controls, and argues that the size of the  $R^2$  change is informative in judging whether the stability of the estimated coefficient of interest is sufficient to argue away the omitted variable problem. Given that the Oster-test is designed for a single key explanatory variable, we focus on the exercises involving Pitch Factor, which we will call the uncontrolled ( $u$ ) and controlled ( $c$ ) regressions, respectively. We denote their estimates and  $R^2$  as  $(\beta_u, R_u^2)$  and  $(\beta_c, R_c^2)$ . Moreover, since the test is designed for linear models, we switch the estimation from Logit to OLS for this analysis.

To obtain an identified set of coefficients, the test strategy relies on assumptions on two more parameters— $\delta$  and  $R_{max}^2$ .  $\delta$  (often referred to as the proportionality parameter) captures the level of selection on unobservables relative to selection on observable controls; a higher  $\delta$  means that the omitted variable problem is more severe.  $R_{max}^2$  is the hypothetical overall  $R^2$  of the model with observables and unobservables. This measure indicates how much of the variation in the outcome variable can be explained by controlling for everything. The bias-adjusted coefficient, denoted as  $\beta_{adj}$  and determined by parameters  $\delta$  and  $R_{max}^2$ , is closely approximated by the equation below (Section 3.2 in Oster (2019)):

$$\beta_{adj} \approx \beta_c - \delta \frac{(\beta_u - \beta_c)(R_{max}^2 - R_c^2)}{R_c^2 - R_u^2}. \quad (3)$$

With this adjusted coefficient  $\beta_{adj}$ , the recommended identified set is the interval between  $\beta_{adj}$  and  $\beta_c$ . We test whether the set safely excludes zero for reasonable parameterizations of  $\delta$  and  $R_{max}^2$ .

**[Insert Table VII Here.]**



In Table VII we report the test results for different combinations of parameters. Table IA.10 in the Appendix shows the raw OLS estimation results used in the Oster test. Following the application of the test in Mian and Sufi (2014), our baseline test takes the values  $R_{max}^2 = \min(2.2R_c^2, 1)$  and  $\delta = 1$ . We show that the adjusted  $\beta$  is close to the estimated value and that we can easily reject the null that  $\beta = 0$ . In fact, the unobservable quality controls appear to be quite unimportant for our estimation. When pushing the  $\delta$  to take a value of 2 and thus implementing the unrestricted estimator in Oster (2019), the identified set is still quite tight at  $[0.019, 0.023]$ . Even when we push the parameterization to very high values— $R_{max}^2 = 1$  and  $\delta = 2$ —we can still reject the null.

This means, in a large set of scenarios, that the effect of omitted quality controls is fairly minimal and that the relation between the pitch features and investment decisions remains robust. In other words, pitch features do not seem to be correlated with funding decisions only because they are proxies for omitted startup quality. We want to acknowledge that it remains an important assumption that unobservables are not more than twice as important as the vast set of observables ( $\delta = 2$ ). Altonji et al. (2005) and Oster (2019) suggest this is appropriate, and the reasoning is that researchers often first focus on the most important set of controls (Angrist and Pischke (2010)). This is also supported by the work on CEO communication styles (Dzieliński, Wagner, and Zeckhauser (2021)), which shows that a CEO’s communication style is at best very weakly related to the fundamentals of the company. Even though  $1 - R_c^2$  means sizable variations are unexplained by the model, this is a shared feature in related literature (Bernstein et al., 2017; Ewens and Townsend, 2020) pointing to the nature of venture investment.

### C. *Performance of Startups*

The evidence so far suggests that the delivery features of a pitch can have independent, robust, and sizable impacts on investors’ decisions. But do these features help investors make better investment decisions? The answer could well be “yes.” For example, entrepreneurs may appear more positive and energetic when their startups are of high quality. Alternatively, the communication and interpersonal skills reflected in the pitch may be valuable assets for a venture. In addition, traits like positivity and warmth might be desirable and advantageous, especially given the challenges and difficulties inherent in entrepreneurship. Under this line of thinking, the result can be interpreted as “a better pitch is a valuable signal for better startups, and therefore these startups receive more funding.”

Building on this reasoning, we examine the long-term performance of startups to further explore the hypothesis. If persuasion features do indeed serve as valuable signals of a startup’s potential, then we would expect startups exhibiting these features to be associated with better future performance. In contrast, if startups that were chosen based on these features underperform other companies conditional on obtaining funding, it could indicate that investors are subject to biases induced by those features. Regardless of the findings, the relation should bear limited causal interpretations. Instead, they offer useful correlations that indicate whether decisions based on these features are associated with better long-term performance.

We perform the analysis using the following model:

$$Performance_i = \alpha + \beta \cdot \text{Pitch Factor}_i + \gamma \cdot \text{Controls}_i + \delta_{FE} + \varepsilon_i. \quad (4)$$

The key explanatory variable is the Pitch Factor. All regressions include pitch content, startup, and team controls. To control for the growth stage, we also include controls for firm age and the squared term of firm age at the time of measurement. Accelerator fixed effects are included to account for variations in investor nurturing and value-adding.

Measuring the performance of early-stage startups is challenging, and we approach this problem in several ways. First, we examine the company’s total employment, which is a standard real-outcome performance measure used in entrepreneurship research ([Puri and Zarutskie, 2012](#); [Adelino, Ma, and Robinson, 2017](#)). Second, we examine whether a startup remains alive and active in development based on its website activities tracked through the Wayback Machine. Third, we examine whether a startup has raised a follow-on round from a VC, typically a Series A round, and the total amount of capital raised from VCs. This measure serves as an interim measure of startup success ([Ewens and Townsend \(2020\)](#)).<sup>24</sup> Finally, we collect information on whether a startup achieves a milestone exit event in the form of an IPO or acquisition. For total employment and amount of VC financing, we use an inverse hyperbolic sine transformation to transform the variables for better empirical properties, and the results remain almost identical to alternative transformation methods such as the log transformation.

---

<sup>24</sup> There are often 12 to 18 months between the accelerator pitch and formal fundraising from VC investors, conditional on the startup’s survival and healthy development. That is, after the accelerator period, there is a natural revelation of quality. Many fail to obtain funding primarily because they do not develop successfully enough to reach the fundraising stage. VC funding outcome measures naturally account for this quality dimension.

[Insert Table VIII Here.]

Table VIII presents the results. In Panel A, we focus on the full sample of our startups, using a similar design as [Huang et al. \(2023\)](#). We estimate a negative and significant coefficient across all the performance measures. Columns (1)–(3) show that startups exhibiting more positive pitch features grow more slowly in employment and are less likely to obtain more VC financing. These measures do not directly address a more nuanced version of startup performance—the right-tail “home-runs.” Column (4) shows that these startups are also less likely to achieve an IPO or acquisition as a milestone exit outcome. This finding is inconsistent with the explanation that features in pitches allow investors to form more accurate beliefs for their investment decisions.

Column (5) uses the Wayback Machine to track the activities of startup websites. Specifically, for each startup website, the Wayback Machine keeps snapshots of the site in the past, at a quarterly frequency or even finer time intervals. We track the changes made to the website over time and aggregate the information to construct a proxy to measure the intensity of website updates. Such a measure can serve as a detailed indicator of a startup’s development or longevity. Conversely, a lack of changes to the website can signal unhealthy development. This metric provides a granular view into startup development/survival. The results are incorporated into column (5) in the performance analysis above, which shows lower website activities in startups with a higher Pitch Factor.

In Panel B, we focus on 270 startups that received seed investment from an early-stage investor (either the incubator program or other angel investors based on Crunchbase and PitchBook databases). The research design in this panel follows the Becker outcome test framework that tracks the performance of the selected group of subjects. This framework is often used in investigating biases in investment decisions in the finance literature ([Butler et al., 2023](#); [Ewens and Townsend, 2020](#); [Duarte et al., 2012](#)). We find consistent results as in Panel A.

Overall, startups with a high Pitch Factor underperform in the long run. We do not interpret this finding to suggest that the ability to deliver a passionate pitch is counter-productive. Instead, our preferred interpretation is that investors are too reluctant to invest in startups with a less positive pitch, and therefore they only do so for the most promising companies. This in turn leads to better performance. In other words, a passionate pitch could lead investors to fund startups that may not merit the funding, suggesting a potential bias.

In Appendix E we provide a detailed conceptual framework to present these potential biases. In this framework, investors fund startups based on a simple cutoff rule, offering funding to all startups

that exceed a certain threshold of predicted success. When investors are biased—either due to a taste-based channel or inaccurate beliefs about the success probability—high-positivity startups may actually underperform. In the case of taste-based bias, the investor continues to derive utility from the startup’s performance but also experiences disutility from investing in startups with low-positivity pitches. Consequently, the investor sets a higher cutoff for them, and as a result, expected performance will be higher for these low-positivity investments, conditional on funding. In the case of inaccurate beliefs, a gap arises between the investor’s perceived performance distribution and the true performance distribution for low-positivity (or high-positivity) startups. These inaccurate beliefs can also lead investors to fund high-positivity startups with greater probability despite the startups having lower (true) expected performance.<sup>25</sup>

#### *D. Heterogeneous Effects Across Gender*

The evidence thus far does not seem to support the explanation that fully rational agents learn from non-content delivery features in persuasion to improve investment decisions. This leaves room for other explanations such as those based on inaccurate beliefs, and/or investor taste and preference. Under these mechanisms, the relation between pitch features and investment choices would vary among subsamples in which those mechanisms may work differently.

One subsample is defined by gender. Women are often judged differently and treated differently in social occasions and economic settings such as hiring or promoting (Bordalo et al. (2019)). For instance, Fredrickson and Roberts (1997) show that women are often more heavily judged on appearance and non-substantive features. Women and men are also expected to follow different gender stereotypes—there are generally higher expectations of men in general ability and task performance domains, while women are expected to be high in terms of warmth, empathy, and altruism (Kite, Deaux, and Haines, 2008; Fiske, 2010; Ellemers, 2018). Meanwhile, competent but less-warm women are biased against, particularly in leading roles (Rudman and Glick, 2001; Eagly and Karau, 2002), such as entrepreneurs. These gender biases govern a wide range of economic activities and outcomes (Goldin and Rouse, 2000; Bagues and Esteve-Volart, 2010; Brooks et al., 2014; Bohren et al., 2019). In recent literature, Sarsons et al. (2021) show that women receive less credit for group work. Cullen and Perez-Truglia (2023) show that employees’ social interactions

---

<sup>25</sup>As in previous literature (most closely related to our work is Ewens and Townsend (2020)), using outcomes to infer biased beliefs requires assumptions on the outcome distributions across groups. That is, our interpretation is based on the assumptions about how the distributions of potential outcomes differ for high- and low-positive startups.

with their managers often favor men, contributing to the gender pay gap. [Biasi and Sarsons \(2022\)](#) show that women are less willing to engage in negotiations over pay, using a setting of public school teachers.

We first separately study startups with male-only or female-only teams (including one-person startups, which for convenience are also called teams). For each team, the metrics are naturally calculated for only people of the same gender and standardized as above but within gender. We apply the same empirical specification as in Eq. (1). By way of comparing the estimation results for the female and male subsamples, we are able to explore whether pitch delivery is more or less relevant for different gendered entrepreneurs.

**[Insert Table IX Here.]**

Table IX columns (1) and (2) present the result. We separately report the regression results for men and women entrepreneurs. Investment decisions on woman-only startups are significantly more sensitive to the performance in the pitch, with coefficients of 0.016 versus 0.218. Column (3) confirms that the magnitudes are statistically different when we perform the analysis using both types of teams and test the coefficients associated with men and women. This result is consistent with the literature on gender stereotyping, which shows that women and men are evaluated differently in social interactions and economic decisions. Not only are women judged more based on appearance and non-substantive features in pitches, but also the sensitivity is in the same direction as the gender stereotype. Investors reward women who fit their stereotypes—that is, those whom they see as warmer and more positive—and aggressively avoid investing in women entrepreneurs who do not fit this profile.

The result cannot be explained by several alternative explanations often discussed in the gender-finance literature. For example, the difference is not explained by different industry compositions of startups founded by male and female entrepreneurs. In our sample, male-only and female-only teams have similar industry distributions over GICS sectors (see Table IA.11). In fact, in this analysis, we control for fixed effects at the industry level. Moreover, this is not due to different probabilities of obtaining funding or different distributional characteristics (e.g., mean, variance) of pitch features among women and men. Additionally, as will be discussed below, this also does not seem to be driven by a simple algorithm-bias explanation in which there are different levels of errors when measuring men and women.

What if a team has both male and female entrepreneurs? In Table IX column (4), we focus on the subsample of startups that have both male and female team members pitching. For each team, we separately calculate the three-V dimensions and the Pitch Factor for women on the team and men on the team. We put those measures jointly in Eq. (1) so we can examine whether the features from women or men carry more weight for the probability that the team will receive funding.

At first glance, men drive the majority of the relations between pitch features and investment decisions in mixed-gender teams. One (very depressing) way to interpret this finding is that women are ignored in the pitches when they co-present with a man—thus, the features of their pitches matter less. Note that this is even though men and women actually speak for similar amounts of time in pitches on average. We acknowledge that the statistical significance of this result is weak, likely due to the small sample size.

Methodological-wise, the divergence of women-men comparisons in single-gender and mixed-gender teams also provides some assurance of the algorithm measurement error problem introduced above. For example, one may worry that it might be more accurate for the algorithm to capture positive emotions from women (Hess et al., 2009; Sun et al., 2019). But in that case, the systematic measurement error would attenuate the results on men in both single-gender and mixed-gender teams, which is not the case. Conversely, if measurement errors are larger for female subjects, the systematic errors would attenuate the results on men in both analyses.

Overall, the evidence suggests that persuasion delivery affects male and female founders differently, in a direction consistent with gender stereotyping and inequality. One further test of this mechanism is to examine the role of reviewers' gender. This is unfortunately not feasible in our setting since investment decisions are often made by groups rather than one individual. We also do not have reliable information on the gender of the lead investor for each startup. It is possible that gender bias is more severe in cross-gender evaluations, while some research shows that gender-related social psychology forces are often salient among both male and female reviewers (Hentschel, Heilman, and Peus (2019)). Similarly, the sample's small number of startups founded by minority racial groups limits any analysis of race.<sup>26</sup>

---

<sup>26</sup> Alternative heterogeneity analysis across the startup sector or founder ages could also be interesting. Our sample does not have enough variations on these dimensions to identify these heterogeneous effects, but they could be useful directions for future studies.

## IV. Experiment: (Inaccurate) Beliefs vs. Taste?

In this last analysis, we explore the economic mechanisms through which pitch delivery affects persuasion effectiveness. We follow [DellaVigna and Gentzkow \(2010\)](#) and explore the mechanisms in two broad categories—taste-based models and miscalibrated/inaccurate beliefs. We want to understand: which of these two mechanisms better explains how the non-content delivery features affect investor decision; and if both exist, how much do they contribute to the investment bias?

### A. A Simple Conceptual Framework

We model venture investors’ investment problems by incorporating the role of persuasion features through two possible channels: a *beliefs channel* and a *taste-based channel*. The beliefs channel works through investors’ expectations: if true, investors would use startup pitch features to form their beliefs about startups’ chances to succeed, accurately or inaccurately. The taste-based channel operates through a standalone component favoring certain pitch features in investors’ preferences: if true, even with the same perceived quality of the startup, investors would still be more likely to invest in certain features.

Formally, an investor  $j$  makes the investment decision on startup  $i$  based on pitch delivery features  $\theta_i$ ; the investor’s beliefs about the success probability of the venture,  $\mu_{ij}$ ; and the precision or confidence level of the belief  $\sigma_{ij}$ . The investment is based on a simple threshold investment rule:

$$I_{ij} = \mathbb{1}_{\{U_{ij} \geq \bar{U}\}}, \quad \text{where } U(\mu_{ij}, \sigma_{ij}, \theta_i) \equiv \gamma_\mu \mu_{ij} + \gamma_\sigma \sigma_{ij} + \kappa \theta_i. \quad (5)$$

In a wide class of decision models,  $\gamma_\mu > 0$ —investors are more willing to make an investment in startups that they believe to have a higher success probability.  $\sigma$  captures the second moment of the belief about the success probability—the larger is  $\sigma$ , the lower is precision and confidence. We should expect  $\gamma_\sigma < 0$  for a risk-averse agent.

The beliefs channel is modeled by allowing  $\mu$  and  $\sigma$  to be determined by hard information about the venture,  $Q_i$ , and the features in the pitch delivery  $\theta_i$ :

$$\mu_{ij} = \lambda_\mu Q_i + \psi_\mu \theta_i, \quad (6a)$$

$$\sigma_{ij} = \lambda_\sigma Q_i + \psi_\sigma \theta_i. \quad (6b)$$



Under this framework,  $\theta_i$  enters the investment decision in two ways—through the impact on beliefs  $\mu$  (the size of the impact is  $\psi_\mu \gamma_\mu \theta_i$ ) and  $\sigma$  (the size of the impact is  $\psi_\sigma \gamma_\sigma \theta_i$ ); and/or through the direct utility gain through the term  $\kappa \theta_i$ . This means the  $\beta$  coefficient in Eq. (1), under our framework, is the combined effect of  $\kappa + \psi_\mu \gamma_\mu + \psi_\sigma \gamma_\sigma$ .

The experiment can help determine whether those channels exist and the relative importance of the channels in driving the main effect. We can expect three potential scenarios:

Scenario	$\psi_{\mu,\sigma}$	$\kappa$	Beliefs Channel	Taste Channel	Decompose $\beta$ in Eq. (1)
1	$\neq 0$	$= 0$	✓	✗	$\beta = \psi_\mu \gamma_\mu + \psi_\sigma \gamma_\sigma$
2	$= 0$	$\neq 0$	✗	✓	$\beta = \kappa$
3	$\neq 0$	$\neq 0$	✓	✓	$\beta = \kappa + \psi_\mu \gamma_\mu + \psi_\sigma \gamma_\sigma$

In these different scenarios, the  $\psi$  captures whether the beliefs channel exists and its strength. The existence and strength of the preference channel hinge on whether  $\kappa > 0$  when we explicitly control for the beliefs of the investor  $\mu$  and  $\sigma$ .

This framework closely follows the mediation analysis framework (see [VanderWeele \(2016\)](#) for a recent survey). The goal of this experimental analysis is to explore the (relative) importance of the potential channels behind the finding that the Pitch Factor affects investment decisions, which fits in mediation analysis. Mediation analysis is “...analyses used to assess the relative magnitude of different pathways and mechanisms by which an exposure may affect an outcome.” ([VanderWeele \(2016\)](#)) In our analysis, model 5 and its empirical form 7 follow closely the traditional analysis framework outlined in [VanderWeele \(2016\)](#).

## B. Experiment Design

Our experiment constructs a setting to allow participants to act as venture investors. The experiment randomly allocates 10 pitch videos to each subject to review, with random ordering. The video pool consists of 62 videos that are highly standardized; they are from the same incubator program and have comparable lengths and resolutions. After viewing each video  $i$ , the subject is asked to answer questions around three main themes: (1) whether she/he would invest in company  $i$ , denoted as  $I_{ij}$ ; (2) her/his expectation of the company’s success probability,  $\mu_{ij}$ , measured between 0 and 100%; and (3) her/his confidence level on her/his decision and expectation,  $\sigma_{ij}$ , measured on a scale of 1 to 5. When eliciting beliefs, the experiment asks both conditional and unconditional

expectations, and the experiment also uses different definitions of success (staying alive, becoming a unicorn, etc.). We obtain the pitch features,  $\theta_{ij}$ , using the same method as in the earlier part of the paper.

The subject pool consists of Master’s students from the Yale School of Management (Yale SOM). All subjects have basic training in core business skills and a basic understanding of entrepreneurial finance. They all completed the Entrepreneurial Finance (MGT 897) class at Yale SOM, with the same instructor, in which they were exposed to startup evaluation in qualitative and quantitative dimensions and VC investments, among other topics. For each subject, we collect basic characteristics (including age, gender, academic background, work experience, ethnicity, etc.). In addition, we elicit her/his unconditional expectations of startup success probability and confidence level before the experiment. The experiment was a bonus assignment in the Yale SOM Entrepreneurial Finance (MGT 897) class. The response rate is 63.75 percent, and 102 subjects (in a class of 160) participated in and successfully finished the experiment, which on average takes 30 minutes.<sup>27</sup> We want to acknowledge that even though MBA students are often used as experimental subjects in research (see, for example, [Lian, Ma, and Wang \(2019\)](#)), they may not be fully representative of all investors in the real world. Nevertheless, given their similar backgrounds, we believe this subject pool is useful in helping us understand underlying economic mechanisms for investor behaviors.

The subject pool is incentivized to participate in the experiment first with a flat bonus grade for the course on a scale of 50, equivalent to 20 percent of the total participation grade. Participants also receive a performance-based pay that is calculated based on the performance of startups they choose: 10 points for each startup that scores in the top 10 percent of performance in either funding amount or total employment among its cohort, an additional 5 points for each startup that stays alive, and –5 points for each startup that fails. The subjects are also incentivized to make an accurate response to the beliefs question—additional investment performance points are added to account for the distance between the realized outcome and the expectation. An example experiment and summary statistics of the subjects and their responses are provided in [Appendix G](#).

---

<sup>27</sup>Out of the 1,020 experimental investment rounds (10 videos  $\times$  102 participants), 68 were incomplete due to an integration glitch between the survey software and video platforms, leaving 952 experimental rounds. Our results are robust to an alternative approach of dropping all subjects with incomplete experimental rounds.

## C. Results

### C.1. Interaction Features and (Inaccurate) Beliefs

We first test the beliefs channel by estimating  $\psi_\mu$  and  $\psi_\sigma$  from Eq. (6a) and (6b). We estimate the model using OLS with fixed effects at the subject  $j$  level. For the belief measures, we use  $\mu$  and  $\sigma$  for three different probabilities—*alive|invested* (“|” means conditional on), *success|invested*, and *alive|notinvested*.

**[Insert Table X Here.]**

Table X shows the relation between beliefs and Pitch Factor  $\theta$ . A more positive Pitch Factor is associated with a higher expectation of the success probability of the startup venture, yet only a mild and statistically weak decrease in the variance—in other words, a weak increase of investor confidence. Regarding the economic magnitude, a one-standard-deviation increase of the Pitch Factor is associated with an increase of 2.3 percentage points in  $P(\text{alive}|\text{invested})$ , which is a 7.4 percent increase from the baseline expectation of 31 percent. These results confirm the existence of a channel through beliefs.

We next quantify whether the relation between Pitch Factor and beliefs is due to accurate or inaccurate belief updating. Note that Table VIII finds that realized performance often negatively correlates with the Pitch Factor. This means that the positive coefficients shown thus far in Table X are a sign of miscalibrated beliefs. The level of inaccuracy relies on the gap between the coefficients on  $\mu$ s in Table X and the realized outcome. We present the relation between realized outcome and the Pitch Factor in column (5). Our key measure of performance is on survival since “success” is difficult to code and maps to the questions asked in the experiment. Consistent with Table VIII, actual survival probability negatively correlates with the Pitch Factor conditional on investment, in contrast to subjects’ beliefs that Pitch Factor would positively predict survival. The miscalibration of beliefs has a magnitude of 0.122 ( $= 0.023 - (-0.099)$ ).

### C.2. Decomposing Inaccurate Beliefs and Preferences

We next estimate the full model, using a logit framework based on Eq. (5):

$$I_{ij} = \underbrace{\kappa \cdot \theta_i}_{\text{Taste}} + \underbrace{\gamma_\mu \cdot \mu_{ij} + \gamma_\sigma \cdot \sigma_{ij}}_{\text{Beliefs}} + \delta_j + \varepsilon_{ij}. \quad (7)$$

Table [XI](#) shows the results. We first confirm in column (1) that, in our experimental sample, the Pitch Factor is positively associated with the probability that the startup company is chosen to be invested. Comparing this experimental estimate with the real-world estimate, the experimental estimate is larger, but this is partially because the video sample used in the experiment (described above) has more higher-quality videos for standardization purposes. After taking this into consideration, economic magnitudes are comparable. Not surprisingly, in columns (2) and (3), when the subject thinks that the company has a higher  $\mu$  or lower  $\sigma$ , the subject is more likely to make an investment.

**[Insert Table [XI](#) Here.]**

Column (4) is the key test of model (7). The Pitch Factor strongly correlates with the investment decision after controlling beliefs, and this suggests that there exists a taste/preference channel through which the pitch features affect investment decisions. In other words, the model supports scenario 3 in Section [IV.A](#) above.

With both channels present, this estimation provides a way to decompose the relative contributions of the two channels to the overall effect of the Pitch Factor. To map the estimated parameters to the framework,  $\kappa = 0.061$  and  $\gamma_\mu = 1.907$ , assuming away the impact of Pitch Factor on  $\sigma$  as supported in Table [X](#). This means that the taste channel leads to an increase (bias) in investment probability by 0.067. The inaccurate beliefs channel leads to an increase in investment by 0.232 ( $= 0.122 \times 1.907$ ). So, the beliefs channel and the taste channel contribute to the bias by 79.2 percent ( $= \frac{0.232}{0.061+0.232}$ ) and 20.8 percent ( $= \frac{0.061}{0.061+0.232}$ ), respectively.<sup>28</sup> This quantitative decomposition may vary from setting to setting, and further explorations of this exercise may be a fruitful path for researchers and practitioners interested in different contexts. However, it is likely that in other financial investment settings, especially those involving professional investors, the inaccurate beliefs channel is more important than the taste-based channel.

## V. Conclusion

It is widely speculated that the delivery of a persuasion matters for the final outcome—sales agents achieve different results selling the same product using the same standard pitch; researchers

---

<sup>28</sup>One caveat is that the 79.2 percent attributed to the beliefs channel could still be an underestimate if solicited beliefs from the experiment are measured with error, attenuating the importance of the belief-based channel. Alternatively, this could be overestimated since the preference-based channel is less salient when the subjects would not interact with the startup founders while the VCs would.

of the same team convince peers to a different level when presenting the same paper using the same slides. Yet there is little evidence on how much and why the delivery features matter, especially in a real-world investment setting.

We shed light on this issue using a novel video-based method applied to a classic setting of persuading investors. We find that non-content delivery features in persuasive interactions have statistically significant and economically sizable effects on investors' decisions. These features do not seem to help investors to make better investment decisions. Instead, our evidence using both archival data and an experiment suggests a bias induced by those features, particularly through leading investors to form inaccurate beliefs.

The results leave many questions unanswered and suggest directions for future research. Conceptually, it will be a fruitful path to explore further the root of the inaccurate beliefs by connecting more closely to behavioral models on persuasion. Among many models, two prominent candidates in persuading investors include categorical and coarse thinking ([Fryer and Jackson, 2008](#); [Mullainathan et al., 2008](#)), and failure to account for repeated information ([DeMarzo, Vayanos, and Zwiebel \(2003\)](#)). We also believe the literature on overconfidence ([Malmendier and Tate \(2005\)](#)), emotions, and affects in behavioral economics could be useful in considering the impact of different pitch styles.

Empirically, our video-based approach is extendable to accommodate more complex settings and measures. The extensions could be along several dimensions. Researchers can track multiple players who sequentially send and receive signals via the three-V dimensions. Moreover, the method can be extended to capture more behaviors—such as gestures, speech fluency, etc. We are hopeful that this paper lays the groundwork for such future research.

## Appendix A: Variable Definitions

Variable	Definition and Construction
<b>A. Visual Metrics</b>	
<i>Visual-Positive</i>	Probability that the facial emotion is happiness by Face++ emotion recognition API
<i>Visual-Negative</i>	Sum of the probabilities that the facial emotion is sadness, anger, fear, and disgust by Face++ emotion recognition API
<i>Visual-Beauty</i>	Beauty scores for the faces in videos by Face++ beauty score API
<b>B. Vocal Metrics</b>	
<i>Vocal-Positive</i>	Probability that the vocal emotion is happiness by the LSTM model in <code>speechemotionrecognition</code>
<i>Vocal-Negative</i>	Probability that the vocal emotion is sadness by the LSTM model in <code>speechemotionrecognition</code>
<i>Vocal-Arousal</i>	Degree of vocal arousal by the SVM model in <code>pyAudioAnalysis</code>
<i>Vocal-Valence</i>	Degree of vocal valence by the SVM model in <code>pyAudioAnalysis</code>
<b>C. Verbal Metrics</b>	
<i>Verbal-Positive</i>	Whether a word is included in the positive category of the LM Master Dictionary (Loughran and McDonald (2011))
<i>Verbal-Negative</i>	Whether a word is included in the negative category of the LM Master Dictionary (Loughran and McDonald (2011))
<i>Verbal-Ability</i>	The direction (−1 or +1) of a word if it is included in the ability category of the NBF dictionary (Nicolas et al. (2019))
<i>Verbal-Warmth</i>	The direction (−1 or +1) of a word if it is included in the warmth category of the NBF dictionary (Nicolas et al. (2019))
<b>D. Startup-level Variables</b>	
<i>I(Invested)</i>	Whether the startup team receives funding from the accelerator
<i>Employment</i>	The inverse hyperbolic sine of the number of employees
<i>Raised VC</i>	Whether the startup team raised another round of VC investment after being funded by the accelerator
<i>VC Amount</i>	The inverse hyperbolic sine of the total amount of raised funding
<i>IPO/Acquisitions</i>	Whether the startup reached a milestone exit event like an IPO or an acquisitions
<i>Website Updates</i>	Annual frequencies of website updates in the three years subsequent to accelerator applications

## References

- Adelino, Manuel, Song Ma, and David Robinson, 2017, Firm age, investment opportunities, and job creation, *Journal of Finance* 72, 999–1038.
- Altonji, Joseph G, Todd E Elder, and Christopher R Taber, 2005, An evaluation of instrumental variable strategies for estimating the effects of catholic schooling, *Journal of Human Resources* 40, 791–821.
- Ambady, Nalini, and Robert Rosenthal, 1992, Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis., *Psychological Bulletin* 111, 256.
- Ambady, Nalini, and Robert Rosenthal, 1993, Half a minute: Predicting teacher evaluations from thin slices of nonverbal behavior and physical attractiveness., *Journal of Personality and Social Psychology* 64, 431.
- Angrist, Joshua D, and Jörn-Steffen Pischke, 2010, The credibility revolution in empirical economics: How better research design is taking the con out of econometrics, *Journal of Economic Perspectives* 24, 3–30.
- Antweiler, Werner, and Murray Z Frank, 2004, Is all that talk just noise? the information content of internet stock message boards, *Journal of Finance* 59, 1259–1294.
- Arkes, Hal R, Lisa Tandy Herren, and Alice M Isen, 1988, The role of potential loss in the influence of affect on risk-taking behavior, *Organizational Behavior and Human Decision Processes* 42, 181–193.
- Åstebro, Thomas, Holger Herz, Ramana Nanda, and Roberto A. Weber, 2014, Seeking the roots of entrepreneurship: Insights from behavioral economics, *Journal of Economic Perspectives* 28, 49–70.
- Bagues, Manuel F, and Berta Esteve-Volart, 2010, Can gender parity break the glass ceiling? evidence from a repeated randomized experiment, *Review of Economic Studies* 77, 1301–1328.
- Barsade, Sigal G, 2002, The ripple effect: Emotional contagion and its influence on group behavior, *Administrative Science Quarterly* 47, 644–675.
- Benjamin, Daniel J, and Jesse M Shapiro, 2009, Thin-slice forecasts of gubernatorial elections, *Review of Economics and Statistics* 91, 523–536.
- Berggren, Niclas, Henrik Jordahl, and Panu Poutvaara, 2010, The looks of a winner: Beauty and electoral success, *Journal of Public Economics* 94, 8–15.
- Bernstein, Shai, Arthur Korteweg, and Kevin Laws, 2017, Attracting early-stage investors: Evidence from a randomized field experiment, *Journal of Finance* 72, 509–538.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman, 2010, What’s advertising content worth? evidence from a consumer credit marketing field experiment, *Quarterly Journal of Economics* 125, 263–306.



- Bestelmeyer, Patricia EG, Sonja A Kotz, and Pascal Belin, 2017, Effects of emotional valence and arousal on the voice perception network, *Social Cognitive and Affective Neuroscience* 12, 1351–1358.
- Biasi, Barbara, and Heather Sarsons, 2022, Flexible wages, bargaining, and the gender gap, *Quarterly Journal of Economics* 137, 215–266.
- Blankespoor, Elizabeth, Bradley E Hendricks, and Gregory S Miller, 2017, Perceptions and price: Evidence from ceo presentations at ipo roadshows, *Journal of Accounting Research* 55, 275–327.
- Bohren, J Aislinn, Kareem Haggag, Alex Imas, and Devin G Pope, 2019, Inaccurate statistical discrimination, *NBER Working Paper* .
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, 2019, Beliefs about gender, *American Economic Review* 109, 739–73.
- Boxell, Levi, 2024, Slanted images: Measuring nonverbal media bias during the 2016 election, *Political Science Research and Methods* forthcoming.
- Braggion, Fabio, Narly Dwarkasing, and Lyndon Moore, 2017, Nothing special about banks: Competition and bank lending in Britain, 1885–1925, *Review of Financial Studies* 30, 3502–3537.
- Brooks, Alison Wood, Laura Huang, Sarah Wood Kearney, and Fiona E Murray, 2014, Investors prefer entrepreneurial ventures pitched by attractive men, *Proceedings of the National Academy of Sciences* 111, 4427–4431.
- Burkhardt, Felix, Astrid Paeschke, Miriam Rolfes, Walter F Sendlmeier, and Benjamin Weiss, 2005, A database of German emotional speech, in *Ninth European Conference on Speech Communication and Technology*.
- Butler, Alexander W, Erik J Mayer, and James P Weston, 2023, Racial disparities in the auto loan market, *Review of Financial Studies* 36, 1–41.
- Choudhury, Prithwiraj, Dan Wang, Natalie A Carlson, and Tarun Khanna, 2019, Machine learning approaches to facial and text analysis: Discovering CEO oral communication styles, *Strategic Management Journal* 40, 1705–1732.
- Clore, Gerald L, Norbert Schwarz, and Michael Conway, 1994, Affective causes and consequences of social information processing., *Handbook of Social Cognition* 1, 323–417.
- Cullen, Zoë, and Ricardo Perez-Truglia, 2023, The old boys’ club: Schmoozing and the gender gap, *American Economic Review* 113, 1703–1740.
- Dahl, Gordon, and Stefano DellaVigna, 2009, Does movie violence increase violent crime?, *Quarterly Journal of Economics* 124, 677–734.
- Dana, Jason, Robyn Dawes, and Nathaniel Peterson, 2013, Belief in the unstructured interview: The persistence of an illusion, *Judgment and Decision making* 8, 512.

- DellaVigna, Stefano, 2009, Psychology and economics: Evidence from the field, *Journal of Economic Literature* 47, 315–72.
- DellaVigna, Stefano, and Matthew Gentzkow, 2010, Persuasion: Empirical evidence, *Annual Review of Economics* 2, 643–669.
- DellaVigna, Stefano, and Ethan Kaplan, 2007, The fox news effect: Media bias and voting, *Quarterly Journal of Economics* 122, 1187–1234.
- DellaVigna, Stefano, and Devin Pope, 2018, What motivates effort? evidence and expert forecasts, *Review of Economic Studies* 85, 1029–1069.
- DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel, 2003, Persuasion bias, social influence, and unidimensional opinions, *Quarterly Journal of Economics* 118, 909–968.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 2018, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* .
- Dewatripont, Mathias, and Jean Tirole, 1999, Advocates, *Journal of Political Economy* 107, 1–39.
- Duarte, Jefferson, Stephan Siegel, and Lance Young, 2012, Trust and credit: The role of appearance in peer-to-peer lending, *Review of Financial Studies* 25, 2455–2484.
- Dzieliński, Michał, Alexander F Wagner, and Richard J Zeckhauser, 2021, Ceo clarity, *Swiss Finance Institute Research Paper* .
- Eagly, Alice H, and Steven J Karau, 2002, Role congruity theory of prejudice toward female leaders., *Psychological Review* 109, 573.
- Ellemers, Naomi, 2018, Gender stereotypes, *Annual Review of Psychology* 69, 275–298.
- Ewens, Michael, 2022, Race and gender in entrepreneurial finance, *Handbook of the Economics of Corporate Finance* Forthcoming.
- Ewens, Michael, and Richard R Townsend, 2020, Are early stage investors biased against women?, *Journal of Financial Economics* 135, 653–677.
- Farrell, Michael, T Clifton Green, Russell Jame, and Stanimir Markov, 2022, The democratization of investment research and the informativeness of retail investor trading, *Journal of Financial Economics* 145, 616–641.
- Fiske, Susan T, 2010, Venus and mars or down to earth: Stereotypes and realities of gender differences, *Perspectives on Psychological Science* 5, 688–692.
- Fiske, Susan T, Amy JC Cuddy, and Peter Glick, 2007, Universal dimensions of social cognition: Warmth and competence, *Trends in cognitive sciences* 11, 77–83.
- Fisman, Raymond, Daniel Paravisini, and Vikrant Vig, 2017, Cultural proximity and loan outcomes, *American Economic Review* 107, 457–92.

- Fredrickson, Barbara L, and Tomi-Ann Roberts, 1997, Objectification theory: Toward understanding women's lived experiences and mental health risks, *Psychology of Women Quarterly* 21, 173–206.
- Fryer, Roland, and Matthew O. Jackson, 2008, A categorical model of cognition and biased decision making, *The B.E. Journal of Theoretical Economics* 8.
- Gandomi, Amir, and Murtaza Haider, 2015, Beyond the hype: Big data concepts, methods, and analytics, *International Journal of Information Management* 35, 137–144.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019, Text as data, *Journal of Economic Literature* 57, 535–74.
- Giannakopoulos, Theodoros, 2015, pyaudioanalysis: An open-source python library for audio signal analysis, *PLoS One* 10, e0144610.
- Goldin, Claudia, and Cecilia Rouse, 2000, Orchestrating impartiality: The impact of "blind" auditions on female musicians, *American Economic Review* 90, 715–741.
- Gómez-Cram, Roberto, and Marco Grotteria, 2022, Real-time price discovery via verbal communication: Method and application to fedspeak, *Journal of Financial Economics* 143, 993–1025.
- Gompers, Paul A, Will Gornall, Steven N Kaplan, and Ilya A Strebulaev, 2020, How do venture capitalists make decisions?, *Journal of Financial Economics* 135, 169–190.
- Gompers, Paul A, and Sophie Q Wang, 2017, And the children shall lead: Gender diversity and performance in venture capital, *NBER Working Paper w23459*.
- Gornall, Will, and Ilya A Strebulaev, 2022, Gender, race, and entrepreneurship: A randomized field experiment on venture capitalists and angels, *Management Science* Forthcoming.
- Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera, 2023, The voice of monetary policy, *American Economic Review* 113, 548–584.
- Gow, Ian D, Steven N Kaplan, David F Larcker, and Anastasia A Zakolyukina, 2016, Ceo personality and firm policies, Technical report, National Bureau of Economic Research.
- Graham, John R, Campbell R Harvey, and Manju Puri, 2016, A corporate beauty contest, *Management Science* 63, 3044–3056.
- Haltiwanger, John, Ron S Jarmin, and Javier Miranda, 2013, Who creates jobs? Small versus large versus young, *Review of Economics and Statistics* 95, 347–361.
- Hamermesh, Daniel S, and Jeff E Biddle, 1994, Beauty and the labor market, *American Economic Review* 1174–1194.
- Hatfield, Elaine, John T Cacioppo, and Richard L Rapson, 1993, Emotional contagion, *Current Directions in Psychological Science* 2, 96–100.
- Hebert, Camille, 2020, Gender stereotypes and entrepreneur financing.

- Hentschel, Tanja, Madeline E Heilman, and Claudia V Peus, 2019, The multiple dimensions of gender stereotypes: a current look at men's and women's characterizations of others and themselves, *Frontiers in Psychology* 10, 11.
- Hess, Ursula, Reginald B Adams, Karl Grammer, and Robert E Kleck, 2009, Face gender and emotion expression: Are angry women more like men?, *Journal of Vision* 9, 19–19.
- Hirshleifer, David, and Tyler Shumway, 2003, Good day sunshine: Stock returns and the weather, *Journal of Finance* 58, 1009–1032.
- Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of Political Economy* 124, 1423–1465.
- Hoberg, Gerard, and Gordon M Phillips, 2010, Product market synergies and competition in mergers and acquisitions: A text-based analysis, *Review of Financial Studies* 23, 3773–3811.
- Hobson, Jessen L, William J Mayew, and Mohan Venkatachalam, 2012, Analyzing speech to detect financial misreporting, *Journal of Accounting Research* 50, 349–392.
- Hochberg, Yael V, 2016, Accelerating entrepreneurs and ecosystems: The seed accelerator model, *Innovation Policy and the Economy* 16, 25–51.
- Howell, Sabrina T, 2019, Reducing information frictions in venture capital: The role of new venture competitions, *Journal of Financial Economics* .
- Huang, Xing, Zoran Ivković, J Jiang, and I Wang, 2023, Angel investment and first impressions, *Journal of Financial Economics* Forthcoming.
- Johnson, Eric J, and Amos Tversky, 1983, Affect, generalization, and the perception of risk., *Journal of personality and social psychology* 45, 20.
- Joo, Jungseock, and Zachary C Steinert-Threlkeld, 2018, Image as data: Automated visual content analysis for political science, *arXiv preprint arXiv:1810.01544* .
- Kamenica, Emir, and Matthew Gentzkow, 2011, Bayesian persuasion, *American Economic Review* 101, 2590–2615.
- Kaplan, Steven N, Mark M Klebanov, and Morten Sorensen, 2012, Which ceo characteristics and abilities matter?, *Journal of Finance* 67, 973–1007.
- Kaplan, Steven N, Berk A Sensoy, and Per Strömberg, 2009, Should investors bet on the jockey or the horse? Evidence from the evolution of firms from early business plans to public companies, *Journal of Finance* 64, 75–115.
- Kaplan, Steven N, and Morten Sorensen, 2021, Are CEOs different?, *Journal of Finance* 76, 1773–1811.
- Kaplan, Steven N, and Per Strömberg, 2004, Characteristics, contracts, and actions: Evidence from venture capitalist analyses, *Journal of Finance* 59, 2177–2210.

- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2021, Measuring technological innovation over the long run, *American Economic Review: Insights* 3, 303–20.
- Kite, Mary E, Kay Deaux, and Elizabeth L Haines, 2008, Gender stereotypes, *Psychology of women: A handbook of issues and theories* 2, 205–236.
- Knapp, Mark L, Judith A Hall, and Terrence G Horgan, 2013, *Nonverbal communication in human interaction* (Cengage Learning).
- Kogan, Shimon, Tobias J Moskowitz, and Marina Niessner, 2021, Social media and financial news manipulation, *Review of Finance* .
- Kortum, Samuel, and Josh Lerner, 2000, Assessing the contribution of venture capital to innovation, *RAND Journal of Economics* 31, 674.
- Krauss, Robert M, William Apple, Nancy Morency, Charlotte Wenzel, and Ward Winton, 1981, Verbal, vocal, and visible factors in judgments of another's affect., *Journal of Personality and Social Psychology* 40, 312.
- Kuhnen, Camelia M, and Brian Knutson, 2011, The influence of affect on beliefs, preferences, and financial decisions, *Journal of Financial and Quantitative Analysis* 46, 605–626.
- Kuhnen, Camelia M, and Alexandra Niessen, 2012, Public opinion and executive compensation, *Management Science* 58, 1249–1272.
- Lerner, Josh, Antoinette Schoar, Stanislav Sokolinski, and Karen Wilson, 2018, The globalization of angel investments: Evidence across countries, *Journal of Financial Economics* 127, 1–20.
- Lian, Chen, Yueran Ma, and Carmen Wang, 2019, Low interest rates and risk-taking: Evidence from individual investment decisions, *Review of Financial Studies* 32, 2107–2148.
- Loewenstein, George F, Elke U Weber, Christopher K Hsee, and Ned Welch, 2001, Risk as feelings., *Psychological Bulletin* 127, 267.
- Loughran, Tim, and Bill McDonald, 2011, When is a liability not a liability? textual analysis, dictionaries, and 10-Ks, *Journal of Finance* 66, 35–65.
- Loughran, Tim, and Bill McDonald, 2016, Textual analysis in accounting and finance: A survey, *Journal of Accounting Research* 54, 1187–1230.
- Malmendier, Ulrike, and Geoffrey Tate, 2005, Ceo overconfidence and corporate investment, *Journal of Finance* 60, 2661–2700.
- Manyika, James, Michael Chui, Brad Brown, Jacques Bughin, Richard Dobbs, Charles Roxburgh, and Angela Hung Byers, 2011, Big data: The next frontier for innovation, competition, *Washington, DC: McKinsey Global Institute* .
- Mayew, William J, and Mohan Venkatachalam, 2012, The power of voice: Managerial affective states and future firm performance, *Journal of Finance* 67, 1–43.

- McNamara, Quinten, Alejandro De La Vega, and Tal Yarkoni, 2017, Developing a comprehensive framework for multimodal feature extraction, in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1567–1574.
- Mehrabian, Albert, 1971, *Silent Messages: Implicit Communication of Emotions and Attitudes*, volume 8 (Wadsworth Belmont, CA).
- Mehrabian, Albert, 1972, *Nonverbal Communication* (Transaction Publishers).
- Mian, Atif, and Amir Sufi, 2014, What explains the 2007–2009 drop in employment?, *Econometrica* 82, 2197–2223.
- Milgrom, Paul, and John Roberts, 1986, Relying on the information of interested parties, *RAND Journal of Economics* 18–32.
- Mobius, Markus M, and Tanya S Rosenblat, 2006, Why beauty matters, *American Economic Review* 96, 222–235.
- Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer, 2008, Coarse thinking and persuasion, *Quarterly Journal of Economics* 123, 577–619.
- Nicolas, Gandalf, Xuechunzi Bai, and Susan T Fiske, 2019, Automated dictionary creation for analyzing text: An illustration from stereotype content .
- Oster, Emily, 2019, Unobservable selection and coefficient stability: Theory and evidence, *Journal of Business & Economic Statistics* 37, 187–204.
- Peng, Lin, Siew Hong Teoh, Yakun Wang, and Jiawen Yan, 2021, Face value: Trait inference, performance characteristics, and market outcomes for financial analysts, *Journal of Accounting Research* Forthcoming.
- Peng, Yilang, 2018, Same candidates, different faces: Uncovering media bias in visual portrayals of presidential candidates with computer vision, *Journal of Communication* 68, 920–941.
- Puri, Manju, and Rebecca Zarutskie, 2012, On the lifecycle dynamics of venture-capital- and non-venture-capital-financed firms, *Journal of Finance* 67, 2247–2293.
- Rosenberg, Shawn W, Lisa Bohan, Patrick McCafferty, and Kevin Harris, 1986, The image and the vote: The effect of candidate presentation on voter preference, *American Journal of Political Science* 108–127.
- Rudman, Laurie A, and Peter Glick, 2001, Prescriptive gender stereotypes and backlash toward agentic women, *Journal of Social Issues* 57, 743–762.
- Sarsons, Heather, Klarita Gërxhani, Ernesto Reuben, and Arthur Schram, 2021, Gender differences in recognition for group work, *Journal of Political Economy* 129, 101–147.
- Schubert, James N, Carmen Strungaru, Margaret Curren, and Wulf Schiefenhovel, 1998, Physische erscheinung und die einschätzung von politischen kandidatinnen und kandidaten, *Biopolitics: Politikwissenschaft jenseits des Kulturismus. Nomos Verlagsgesellschaft, Baden-Baden* .

- Schwartzstein, Joshua, and Adi Sunderam, 2021, Using models to persuade, *American Economic Review* 111, 276–323.
- Shorrocks, Anthony F, 2013, Decomposition procedures for distributional analysis: a unified framework based on the shapley value, *Journal of Economic Inequality* 11, 99–126.
- Stigler, George J, 1961, The economics of information, *Journal of Political Economy* 69, 213–225.
- Strahan, Carole, and Donald G Zytowski, 1976, Impact of visual, vocal, and lexical cues on judgments of counselor qualities., *Journal of Counseling Psychology* 23, 387.
- Sun, Tony, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang, 2019, Mitigating gender bias in natural language processing: Literature review, *arXiv preprint arXiv:1906.08976* .
- Tausczik, Yla R, and James W Pennebaker, 2010, The psychological meaning of words: Liwc and computerized text analysis methods, *Journal of Language and Social Psychology* 29, 24–54.
- Tetlock, Paul C, 2007, Giving content to investor sentiment: The role of media in the stock market, *Journal of Finance* 62, 1139–1168.
- Tetlock, Paul C, Maytal Saar-Tsechansky, and Sofus Macskassy, 2008, More than words: Quantifying language to measure firms’ fundamentals, *Journal of Finance* 63, 1437–1467.
- Todorov, Alexander, Anesu N Mandisodza, Amir Goren, and Crystal C Hall, 2005, Inferences of competence from faces predict election outcomes, *Science* 308, 1623–1626.
- VanderWeele, Tyler J, 2016, Mediation analysis: a practitioner’s guide, *Annual Review of Public Health* 37, 17–32.






Figure 1. Data structure and processing procedure. This figure illustrates the flow of processing video data using a real video example.

**Table I**  
**Summary Statistics of Videos and Startups**

This table provides descriptive statistics of pitch videos and the underlying startups in our sample. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles. Panel A reports basic information of the pitch videos. Panel B reports the characteristics and long-term performance of startups measured as of July 2023 from Crunchbase and PitchBook. Panel C reports the summary statistics of the startup teams based on the presenting team members. Team member background information is collected from LinkedIn. Panel D reports the summary statistics of informational content controls.

Panel A: Summary Statistics of Video Pitches						
	N	Mean	STD	25%	50%	75%
Duration (second)	1,139	83.43	39.51	60.00	68.00	97.00
Video Size (MB)	1,139	18.16	18.37	6.88	12.86	22.64
Number of Words	1,139	228.53	107.76	163.00	201.00	262.00
Number of Sentences	1,139	15.91	7.33	11.00	14.00	19.00
Number of Views (YouTube)	1,139	764.37	6956.02	31.00	79.00	197.00
Number of Likes (YouTube)	1,139	1.51	6.60	0.00	0.00	1.00
Number of Dislikes (YouTube)	1,139	0.15	0.65	0.00	0.00	0.00
Panel B: Summary Statistics of Startups (as of July 2023)						
	N	Mean	STD	25%	50%	75%
Invested by Accelerator	1,139	0.09	0.28	0.00	0.00	0.00
Firm Age	1,139	6.20	2.01	5.00	6.00	8.00
<i>Full Sample</i>						
Number of Employees	1,139	7.53	43.88	0.00	0.00	5.00
Raised VC	1,139	0.10	0.30	0.00	0.00	0.00
Total Funding Amount (\$000)	1,139	1388.14	20004.70	0.00	0.00	0.00
IPO/Acquisitions	1,139	0.03	0.16	0.00	0.00	0.00
Website Updates	1,139	0.80	1.33	0.00	0.00	1.33
<i>Conditional on Seed Funding</i>						
Number of Employees	270	22.41	74.32	5.00	5.00	30.00
Raised VC	270	0.40	0.49	0.00	0.00	1.00
Total Funding Amount (\$000)	270	5854.48	40825.78	0.00	0.00	450.00
IPO/Acquisitions	270	0.09	0.29	0.00	0.00	0.00
Website Updates	270	1.18	1.39	0.00	0.64	2.12

Panel C: Summary Statistics of Teams

	N	Mean	STD	25%	50%	75%
Number of People	1,139	1.74	0.84	1.00	2.00	2.00
Single-Member	1,139	0.46	0.50	0.00	0.00	1.00
Multi-Member	1,139	0.54	0.50	0.00	1.00	1.00
Men-Only	1,139	0.49	0.50	0.00	0.00	1.00
Women-Only	1,139	0.27	0.45	0.00	0.00	1.00
Mixed Gender	1,139	0.24	0.43	0.00	0.00	0.00
Has LinkedIn	1,139	0.61	0.49	0.00	1.00	1.00
Prior Senior Position	693	0.76	0.42	1.00	1.00	1.00
Prior Startup Experience	693	0.49	0.50	0.00	0.00	1.00
Elite University	693	0.10	0.30	0.00	0.00	0.00
Master Degree	693	0.32	0.47	0.00	0.00	1.00
PhD Degree	693	0.05	0.22	0.00	0.00	0.00

Panel D: Summary Statistics of Informational Content Controls

	N	Mean	STD	25%	50%	75%
<i>Textual Similarity</i>						
Idea Novelty (PB)	1,139	1.06	0.03	1.03	1.05	1.07
Idea Novelty (10K)	1,139	1.09	0.23	1.05	1.07	1.11
<i>Dictionary-based</i>						
Concrete Number	1,139	0.61	0.49	0.00	1.00	1.00
Cash Flow	1,139	0.16	0.37	0.00	0.00	0.00
Competition	1,139	0.06	0.23	0.00	0.00	0.00
Employment	1,139	0.08	0.28	0.00	0.00	0.00
Readiness	1,139	0.20	0.40	0.00	0.00	0.00
Technology	1,139	0.25	0.43	0.00	0.00	0.00
Data AI	1,139	0.42	0.49	0.00	0.00	1.00
<i>LIWC</i>						
Focus Past	1,139	2.54	1.90	1.14	2.21	3.64
Focus Present	1,139	11.63	3.41	9.35	11.43	13.53
Focus Future	1,139	1.28	1.12	0.53	1.09	1.85
Concreteness	1,139	0.71	3.07	-1.26	0.69	2.76
Informal	1,139	0.70	1.77	0.00	0.47	0.95

**Table II**  
**Summary Statistics of Pitching Behavior Metrics**

This table provides summary statistics of the pitch delivery features. In Panel A, for each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles. Variables are categorized into vocal, video, and verbal. The last two columns in Panel A report the factor loading and uniqueness of each feature when performing the principal component factor analysis to generate the single Pitch Factor that captures the maximum variance in the set of pitch features (Visual-Beauty is excluded as it is a static appearance feature). Panel B provides correlations of the features extracted from the pitches. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

**Panel A: Summary Statistics of Unstandardized Features**

	N	Mean	STD	25%	50%	75%	Pitch Factor Loading	Uniqueness
<i>Visual (Facial)</i>								
Visual-Positive	1,139	0.17	0.16	0.05	0.12	0.25	0.08	0.29
Visual-Negative	1,139	0.15	0.14	0.06	0.11	0.20	−0.14	0.27
Visual-Beauty	1,139	0.58	0.08	0.54	0.59	0.64		
<i>Vocal (Audio)</i>								
Vocal-Positive	1,139	0.09	0.05	0.06	0.08	0.12	0.39	0.41
Vocal-Negative	1,139	0.01	0.01	0.01	0.01	0.02	−0.30	0.43
Vocal-Arousal	1,139	0.55	0.27	0.39	0.58	0.76	0.91	0.15
Vocal-Valence	1,139	0.44	0.22	0.31	0.46	0.59	0.88	0.18
<i>Verbal (Text)</i>								
Verbal-Positive	1,139	0.01	0.01	0.01	0.01	0.02	0.03	0.35
Verbal-Negative	1,139	0.01	0.01	0.00	0.01	0.01	−0.14	0.42
Verbal-Warmth	1,139	0.02	0.01	0.01	0.01	0.02	0.06	0.62
Verbal-Ability	1,139	0.03	0.02	0.02	0.03	0.05	0.06	0.56

**Panel B: Correlations of the Features**

	(1)	(2)	(3)	(4)	(5)
(1) Visual-Positive	1.00				
(2) Visual-Negative	−0.12***	1.00			
(3) Visual-Beauty	−0.02	−0.20***	1.00		
(4) Vocal-Positive	0.16***	0.07**	−0.05*	1.00	
(5) Vocal-Negative	0.05*	0.06**	0.01	−0.07**	1.00
(6) Vocal-Arousal	0.02	−0.07**	0.05*	0.24***	−0.15***
(7) Vocal-Valence	−0.02	−0.07**	0.09***	0.13***	−0.12***
(8) Verbal-Positive	0.01	0.03	−0.01	0.02	−0.06*
(9) Verbal-Negative	−0.10***	0.04	−0.01	0.02	−0.04
(10) Verbal-Warmth	−0.05	0.00	0.01	−0.01	−0.01
(11) Verbal-Ability	0.00	0.02	−0.02	0.04	0.02
<i>Continued</i>	(6)	(7)	(8)	(9)	(10)
(6) Vocal-Arousal	1.00				
(7) Vocal-Valence	0.75***	1.00			
(8) Verbal-Positive	−0.01	0.01	1.00		
(9) Verbal-Negative	−0.08***	−0.07**	0.00	1.00	
(10) Verbal-Warmth	0.02	0.04	0.03	−0.05*	1.00
(11) Verbal-Ability	0.01	0.04	0.08***	−0.03	−0.02

**Table III**  
**Pitch Factor and Investment Decisions**

Logit regressions, marginal effect. The analysis is obtained using the following model:

$$I(Invested) = \alpha + \beta \cdot \text{Pitch Factor} + \gamma \cdot \text{Controls} + \delta_{FE} + \varepsilon.$$

$I(Invested)$  takes a value of one if the startup team was chosen by the accelerator, and zero otherwise. The Pitch Factor is standardized into a zero-mean variable with a standard deviation of one. Controls also include measures capturing the informational content of pitches, including ideal novelty, dictionary-based measures of content, and linguistic characteristics of pitches. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
	Dependent Var: $I(invested)$				
Pitch Factor	0.030*** (0.007)	0.030*** (0.007)	0.027*** (0.007)	0.029*** (0.007)	0.027*** (0.007)
<i>Textual Similarity</i>					
Idea Novelty (PB)		0.705*** (0.179)			0.625*** (0.219)
Idea Novelty (10K)		0.037** (0.016)			0.048*** (0.017)
<i>Dictionary-based</i>					
Concrete Number			0.039* (0.020)		0.031 (0.019)
Cash Flow			0.050** (0.023)		0.055*** (0.020)
Competition			0.056** (0.025)		0.049** (0.024)
Employment			−0.041 (0.037)		−0.028 (0.034)
Readiness			−0.012 (0.022)		−0.008 (0.021)
Technology			0.013 (0.024)		0.014 (0.024)
Data AI			0.012 (0.015)		0.011 (0.014)
<i>LIWC</i>					
Focus Past				0.009*** (0.003)	0.007** (0.003)
Focus Present				−0.005*** (0.002)	−0.005** (0.002)
Focus Future				−0.010* (0.005)	−0.007 (0.007)
Concreteness				0.007*** (0.002)	0.007*** (0.002)
Informal				0.003 (0.002)	0.003* (0.002)
Observations	1,139	1,139	1,139	1,139	1,139
Pseudo $R^2$	0.193	0.212	0.225	0.228	0.273
Accelerator FE	Y	Y	Y	Y	Y

**Table IV**  
**Features in Pitch Delivery and Investment Decisions**

Logit regressions, marginal effect,  $N = 1,139$ . The analysis is obtained using the following model:

$$I(Invested) = \alpha + \beta \cdot \text{Pitch Factor/Features} + \gamma \cdot \text{Controls} + \delta_{FE} + \varepsilon.$$

$I(Invested)$  takes a value of one if the startup team was chosen by the accelerator, and zero otherwise. All pitch feature variables are standardized into a zero-mean variable with a standard deviation of one. All regressions include pitch content controls and accelerator FE. Startup/team control variables include founders' education background (whether they have a master's degree or a PhD degree, whether they attended an elite university, defined as the U.S. News & World Report's Top 10), founders' prior work experience (whether they have prior entrepreneurship experience, whether they ever held a senior position in prior employment), team size, and video resolution. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	Dependent Var: $I(Invested)$			w/ Content Controls			w/ Content and Startup/Team Controls		
	Marginal Effect	S.E.	Pseudo $R^2$	Marginal Effect	S.E.	Pseudo $R^2$	Marginal Effect	S.E.	Pseudo $R^2$
Pitch-Factor	0.027***	(0.007)	0.273	0.024***	(0.007)	0.325			
<i>Visual (Facial)</i>									
Visual-Positive	0.015***	(0.005)	0.260	0.013**	(0.006)	0.314			
Visual-Negative	-0.029***	(0.007)	0.272	-0.030***	(0.007)	0.329			
Visual-Beauty	0.014**	(0.006)	0.259	0.012*	(0.006)	0.313			
<i>Vocal (Audio)</i>									
Vocal-Positive	0.008**	(0.004)	0.256	0.010**	(0.005)	0.312			
Vocal-Negative	-0.036**	(0.016)	0.261	-0.039**	(0.015)	0.317			
Vocal-Arousal	0.021**	(0.009)	0.265	0.017**	(0.008)	0.317			
Vocal-Valence	0.021***	(0.007)	0.265	0.018***	(0.007)	0.317			
<i>Verbal (Text)</i>									
Verbal-Positive	-0.008	(0.010)	0.255	-0.009	(0.010)	0.311			
Verbal-Negative	-0.028***	(0.007)	0.271	-0.026***	(0.007)	0.324			
Verbal-Warmth	0.029***	(0.009)	0.276	0.030***	(0.007)	0.333			
Verbal-Ability	-0.045***	(0.007)	0.315	-0.040***	(0.006)	0.361			

**Table V**  
**Sample Selection of Available Videos**

Logit regressions, marginal effect. This table investigates the sample selection issue of the video sample. The analysis restricts to videos that were uploaded between 2018 and July 2019. By the end of March 2020, 126 videos, or 23.9 percent, were selected out (unlisted, privatized, or removed) from the hosting platforms. The analysis investigates the relation between a video being “selected out” (made private, unlisted, or completely removed) and pitch delivery features and the outcomes of the startup. We include startup/team control variables and the set of content control variables. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Video Selected Out = 1					
Pitch Factor	0.006 (0.022)	0.017 (0.021)				
I( <i>Invested</i> )			−0.042 (0.183)	−0.043 (0.171)		
VC Invested					−0.011 (0.064)	−0.034 (0.061)
Observations	527	527	527	527	527	527
Pseudo $R^2$	0.000	0.053	0.000	0.052	0.000	0.053
Startup/Team Controls		Y		Y		Y
Content Controls		Y		Y		Y
Accelerator FE		Y		Y		Y



**Table VI**  
**Measure Construction—Full Video and Full Channels**

Logit regressions, marginal effect. Panel A uses horse-race regressions to compare Pitch Factor (constructed using complete videos) with Pitch Factor First Slice (constructed using the slice of the first word time interval in each video) and Pitch Factor Random Slice (constructed using the slice of a random word time interval in each video). Panel B uses horse-race regressions to compare Pitch Factor (constructed from three-V channels jointly) and Visual/Vocal/Verbal Factor (constructed from three-V channels separately).  $I(Invested)$  takes a value of one if the startup team was chosen by the accelerator, and zero otherwise. All pitch feature variables are standardized into a zero-mean variable with a standard deviation of one. Column (6) reports the percentage  $R^2$  contributions of each variable from a Shapley-Owen decomposition, using the same specification as in column (5). We include startup/team control variables and the set of content control variables. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Full Video and Thin Slice						
	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent Var: I( <i>Invested</i> )					<i>R</i> <sup>2</sup> Contribution
Pitch Factor		0.025*** (0.007)		0.027** (0.012)	0.026** (0.011)	66.55%
Pitch Factor (First Slice)	0.016* (0.008)	0.002 (0.008)			0.003 (0.008)	16.78%
Pitch Factor (Random Slice)			0.020*** (0.006)	-0.002 (0.010)	-0.002 (0.010)	16.66%
Observations	1,139	1,139	1,139	1,139	1,139	
Pseudo <i>R</i> <sup>2</sup>	0.290	0.302	0.294	0.302	0.302	
Startup/Team Controls	Y	Y	Y	Y	Y	
Content Controls	Y	Y	Y	Y	Y	
Accelerator FE	Y	Y	Y	Y	Y	

Panel B: Full Channels and Individual Channel						
	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent Var: I( <i>Invested</i> )					<i>R</i> <sup>2</sup> Contribution
Pitch Factor					0.053** (0.024)	39.44%
Vocal Factor	0.022*** (0.008)			0.023*** (0.007)	-0.029 (0.026)	23.46%
Visual Factor		0.025*** (0.006)		0.025*** (0.006)	0.020*** (0.006)	34.07%
Verbal Factor			0.003 (0.008)	0.000 (0.008)	-0.007 (0.009)	3.04%
Observations	1,139	1,139	1,139	1,139	1,139	
Pseudo <i>R</i> <sup>2</sup>	0.296	0.301	0.285	0.263	0.317	
Startup/Team Controls	Y	Y	Y	Y	Y	
Content Controls	Y	Y	Y	Y	Y	
Accelerator FE	Y	Y	Y	Y	Y	

**Table VII**  
**Features in Pitches and Investment Decisions—Oster Test**

This table tests the role of omitted and unobservable control variables in explaining the relation between the Pitch Factor and the venture investment decision, using the test designed in Oster (2019). To implement, we estimate a linear model of

$$I(Invested) = \alpha + \beta \cdot \text{Pitch Factor} + \gamma \cdot \text{Controls} + \delta_{FE} + \varepsilon.$$

first without any control variables through which we obtain  $\beta_u$  and  $R_u^2$ , and then with the added startup/team control variable, through which we obtain  $\beta_c$  and  $R_c^2$ . The set of startup/team control variables is identical to that in Table IV. The raw OLS estimates used in this test are provided in Appendix Table IA.10.

For any given test parameter combination  $\delta$  and  $R_{max}^2$ , Oster (2019) defines the bias-adjusted coefficient, denoted as  $\beta_{adj}$  that is determined by parameters  $\delta$  and  $R_{max}^2$ , to be closely approximated by (strictly equal to when  $\delta = 1$ )

$$\beta_{adj} \approx \beta_c - \delta \frac{(\beta_u - \beta_c)(R_{max}^2 - R_c^2)}{R_c^2 - R_u^2}.$$

With this adjusted coefficient  $\beta_{adj}$ , the recommended identified set is the interval between  $\beta_{adj}$  and  $\beta_c$ . In the table, we report the adjusted  $\beta$  and identified set for different combinations of parameters, and we also report whether the identified set rejects the null of  $\beta = 0$  and the  $\delta$  value to make certain  $R_{max}^2$  reach zero.

$R_{max}^2 = \min(2.2R_c^2, 1)$						
$\delta = 1$			$\delta = 2$			$\delta$ s.t. $\beta_{adj} = 0$
$\beta_{adj}$	Identified Set	Reject Null?	$\beta_{adj}$	Identified Set	Reject Null?	
0.021	[0.021,0.023]	Y	0.019	[0.019,0.023]	Y	8.062

$R_{max}^2 = \min(3R_c^2, 1)$						
$\delta = 1$			$\delta = 2$			$\delta$ s.t. $\beta_{adj} = 0$
$\beta_{adj}$	Identified Set	Reject Null?	$\beta_{adj}$	Identified Set	Reject Null?	
0.020	[0.020,0.023]	Y	0.016	[0.016,0.023]	Y	4.890

$R_{max}^2 = 1$						
$\delta = 1$			$\delta = 2$			$\delta$ s.t. $\beta_{adj} = 0$
$\beta_{adj}$	Identified Set	Reject Null?	$\beta_{adj}$	Identified Set	Reject Null?	
0.016	[0.016,0.023]	Y	0.006	[0.006,0.023]	Y	2.482

**Table VIII**  
**Features in Pitches and Long-Term Performance of Startups**

OLS regressions (columns (1) (3) (5)) and Logit regressions, marginal effect (columns (2) and (4)). The analysis is obtained using the following model conditional on receiving funding from a VC:

$$Performance = \alpha + \beta \cdot \text{Pitch Factor} + \gamma \cdot \text{Controls} + \delta_{FE} + \varepsilon.$$

Employment is the inverse hyperbolic sine of number of employees. Raised VC is a dummy variable that takes a value of one if a startup raised another round of VC investment after receiving seed funding. VC Amount is the inverse hyperbolic sine of total amount of VC investment that a startup has raised. IPO/Acquisitions indicates whether the startup reached a milestone exit through an IPO or an acquisition. Website Update is a dummy variable capturing the startup's website update intensity as tracked by WayBack Machine. Panel A uses all startups in our sample and controls for a dummy variable of whether a startup receives seed funding; Panel B uses only the sample of startups that obtained funding from seed investors. All performance variables are as of August 2023. Pitch Factor is standardized into a zero-mean variable with a standard deviation of one. Standard errors clustered at the industry-level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Full Sample ( $N = 1,139$ )					
	(1) Employment	(2) Raised VC	(3) VC Amount	(4) IPO/Acquisitions	(5) Website Update
Pitch Factor	-0.071*** (0.018)	-0.009*** (0.003)	-0.041** (0.017)	-0.011*** (0.002)	-0.079* (0.042)
(Pseudo) $R^2$	0.295	0.589	0.258	0.408	0.203
Age Controls	Y	Y	Y	Y	Y
Content Controls	Y	Y	Y	Y	Y
Startup/Team Controls	Y	Y	Y	Y	Y
Accelerator FE	Y	Y	Y	Y	Y
Region FE	Y	Y	Y	Y	Y

Panel B: Invested Subsample ( $N = 270$ )					
	(1) Employment	(2) Raised VC	(3) VC Amount	(4) IPO/Acquisitions	(5) Website Update
Pitch Factor	-0.284*** (0.070)	-0.033** (0.014)	-0.201** (0.089)	-0.038*** (0.014)	-0.298** (0.092)
(Pseudo) $R^2$	0.195	0.199	0.188	0.189	0.426
Age Controls	Y	Y	Y	Y	Y
Content Controls	Y	Y	Y	Y	Y
Startup/Team Controls	Y	Y	Y	Y	Y
Accelerator FE	Y	Y	Y	Y	Y
Region FE	Y	Y	Y	Y	Y

**Table IX**  
**Gender Differences in the Pitch-Investment Relation**

Logit regressions, marginal effect. The analysis is obtained using the following model:

$$I(Invested) = \alpha + \beta \cdot \text{Pitch Factor/Features} + \gamma \cdot \text{Controls} + \delta_{FE} + \varepsilon.$$

$I(Invested)$  takes value of one if the startup team was chosen by the accelerator, and zero otherwise. Pitch Factor is standardized into a zero-mean variable with a standard deviation of one. Columns (1) and (2) separately analyze for men-only and women-only startup teams and Pitch Factor is standardized separately. Column (3) pools the sample of men-only and women-only startup teams. Pitch Factor (Men) is Pitch Factor times Pure Men Team dummy. Pitch Factor (Women) is Pitch Factor times Pure Women Team dummy. Column (4) analyzes mixed-gender teams, but Pitch Factors are calculated for men and women separately in each video. We also include the set of startup/team control variables and the set of content control variables. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
	Dependent Var: $I(Invested)$			
	Single-Gender Teams		Mixed-Gender Teams	
	Men	Women	Pooled	Pooled
Pitch Factor (Men)	0.016* (0.009)		0.016* (0.009)	0.065*** (0.022)
Pitch Factor (Women)		0.218*** (0.061)	0.074*** (0.028)	-0.002 (0.025)
$p$ -value of Men vs. Women Test			0.058* (0.028)	0.114
Observations	559	310	869	270
Pseudo $R^2$	0.235	0.662	0.268	0.740
Startup/Team Controls	Y	Y	Y	Y
Content Controls	Y	Y	Y	Y
Accelerator FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y

**Table X**  
**Experiment Results: Pitch Factor and Investor Beliefs**

OLS regressions. This table investigates the relation between the Pitch Factor and investor beliefs in an experiment setting (columns (1) to (4)) and realized startup performance conditional on being invested (column (5)).  $P(\text{alive}|\text{invested})$   $\mu$  and  $P(\text{alive}|\text{invested})$   $\sigma$  are subjects' beliefs and precision of beliefs on the probability of a startup to be alive three years later conditional on raising funding.  $P(\text{success}|\text{invested})$   $\mu$  and  $P(\text{success}|\text{invested})$   $\sigma$  are subjects' beliefs and precision of beliefs on the probability that a startup will be a success conditional on raising funding. In column (5),  $\text{alive}|\text{invested}$  is a dummy variable which takes a value of one if a startup's website is still operational as of July 2019. The sample focuses on those startups that obtained venture funding, i.e., conditional on a startup being invested. The Pitch Factor is standardized into a zero-mean variable with a standard deviation of one. All analysis controls for subject individual fixed effects and controls of startup/team and content. Standard errors two-way clustered at the startup and subject levels are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1) $P(\text{alive} \text{invested})$ $\mu$	(2) $P(\text{alive} \text{invested})$ $\sigma$	(3) $P(\text{success} \text{invested})$ $\mu$	(4) $P(\text{success} \text{invested})$ $\sigma$	(5) $\text{alive} \text{invested}$ Realized
Pitch Factor ( $\theta$ )	0.023** (0.010)	-0.022 (0.028)	0.019** (0.007)	-0.017 (0.032)	-0.099* (0.055)
Observations	952	952	952	952	495
$R^2$	0.575	0.548	0.570	0.524	0.743
Startup/Team Controls	Y	Y	Y	Y	Y
Content Controls	Y	Y	Y	Y	Y
Subject FE	Y	Y	Y	Y	Y

**Table XI**  
**Experiment Results: Inaccurate Beliefs, Tastes, and Investment**

Logit regressions, marginal effect. This table estimates the model of investment decisions in an experiment setting, as in Eq. (7).  $I(Invested)$  takes a value of one if a subject decides to invest in a startup team in the experiment, and zero otherwise.  $\mu(alive|invested)$  and  $\sigma(alive|invested)$  are subjects' beliefs and precision of beliefs on the probability that a startup will be alive three years later conditional on receiving funding. The Pitch Factor is standardized into a zero-mean variable with a standard deviation of one. All analysis controls for subject individual fixed effects and control of startup/team characteristics and pitch content. Standard errors two-way clustered at the startup and subject levels are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
	Dependent Var: $I(Invested)$			
Pitch Factor ( $\theta$ )	0.100*** (0.035)			0.061** (0.025)
$\mu(alive invested)$		1.983*** (0.132)		1.907*** (0.139)
$\sigma(alive invested)$			-0.121*** (0.041)	-0.047 (0.031)
Observations	952	952	952	952
Pseudo $R^2$	0.148	0.428	0.146	0.441
Startup/Team Controls	Y	Y	Y	Y
Content Controls	Y	Y	Y	Y
Subject FE	Y	Y	Y	Y

**Internet Appendix for**  
**Persuading Investors: A Video-Based Study**

ALLEN HU    and    SONG MA

---

\* Citation format: Hu, Allen and Song Ma, Internet Appendix for “Persuading Investors: A Video-Based Study,” *The Journal of Finance*, DOI: XXXXX. Please note: Wiley is not responsible for the content or functionality of any additional information provided by the authors. Any queries (other than missing materials) should be directed to the authors of the article.



## A. Collecting Video Data and Startup/Team Information

### A.1. Collecting Videos from Video Platforms

When startups apply to accelerator programs, they are required (or highly recommended) to record and submit a standardized self-introductory pitch video as part of the application process. Figure IA.1 shows such examples from those accelerators’ application systems. These videos, rather than being submitted to the accelerators directly, are submitted through uploading to a public multimedia platform, such as YouTube or Vimeo, and then providing the url links to these videos in application forms.

We use an automatic searching script for two public video-sharing websites, YouTube and Vimeo. Integrated with query APIs, our web crawler returns a list of video indices according to a set of predefined keywords, which include but are not limited to the names of these accelerators, “startup accelerator application video”, “accelerator application videos” and so on. We first obtain the full list of potential videos returned by each keyword search (there is a limit of returned videos by YouTube), and then filter the potential videos by a combination of different conditions on video info obtained along with the video itself. Filtering variables include but are not limit to data format, duration, title, and annotation.

**Table IA.1**  
**List of Searching Keywords for Collecting Videos**

This table shows the list of keywords we use for searching and collecting the pitch videos from Youtube and Vimeo. The *YEAR* takes values from 2005 to 2019.

Keywords
YC Application Videos
Y Combinator Application Videos
MassChallenge Application Videos
500 Startups Application Videos
Techstars Application Videos
AngelPad Application Videos
Y Combinator Application Videos + <i>YEAR</i>
Techstars Application Videos + <i>YEAR</i>
500 Startups Application Videos + <i>YEAR</i>
AngelPad Application Videos + <i>YEAR</i>

We also employ startup names listed on accelerators’ web pages to expand our video data set.

Specifically, we first obtain the full list of startups accelerated by the accelerator each year if such a list is published on the accelerator's website. Then our script automatically searches these startup names and checks the first three results returned by the search API. A match is defined as having both the startup name and the accelerator name appear in the video title or annotation.

It is worth noting that if one company has more than one video in our sample, we only keep the video recorded first. There are 33 such firms in our analysis, which make up only 2.90% of our sample. These firms have multiple videos because of the following reasons. First, there are some entrepreneur teams applying to different accelerators. Second, there are some teams that applied to the same accelerator multiple times. For these firms, we only keep their videos and outcomes in the first application.

In total we obtain 1,139 videos. Table [IA.2](#) describes the sample, in which the number of videos is reported by accelerator (Panel A) and by year (Panel B). Y Combinator contributes the largest number of application videos, followed by MassChallenge and Techstars. Among all the companies that applied, 97 (8.52%) were chosen by the accelerator program, and 248 (21.77%) were invested by any venture investor (accelerator or angels/VCs). The videos are more available for recent years due to the increase in video requirements in the application.

After collecting the videos, we parse each video web page to collect other relevant information. This includes the video's duration, upload date, title, annotation, subtitle, and uploader ID. This set of information also allows us to identify the startup almost perfectly. Specifically, by scrutinizing video titles and annotations, we double-check names of the startups and names of the accelerators they are applying for. If the startup name cannot be identified from these items, we search the uploader name on LinkedIn and back out the company information. It is common that many people have the same name on LinkedIn, so to verify that the person on LinkedIn is the founder, we also double-check the name, background, experience, and even photos.

### Panel A: Y Combinator

**FOUNDERS**

Please provide the email addresses of the other cofounders in the startup. No need to add yours again. **Founders must have at least 10% equity in the company.** We will send an email to each founder to fill out additional information about themselves.

Please enter the url of a 1 minute unlisted (not private) YouTube or Youku video introducing the founder(s). This video is an important part of the application. ([Follow the Video Guidelines.](#))

http://

How many founders are on the team?  
(Fill out this number of founder profiles)

1

### Panel B: MassChallenge

Video elevator pitch url

Upload your 1-3 minute video pitch to Vimeo or Youtube. Paste the shared link here.

**LOCATION**

The main office or headquarters of your company. If your company does not have an address, use your home address

**Country \***

Select country

**State/Region/Province**

---

**City**

### Panel C: 500 Startups

What channel(s) or tool(s) are fueling your customer growth?\*

Please upload your latest capitalization table here:\*

Choose File No file chosen

Please provide a link to your pitch deck here:\*

Please check the sharing settings. If your deck is accessible in incognito, you are good to go.

Anything else you want to tell us?

How did you hear about us?\*

- Please Select -

### Panel D: Techstars

Twitter

Facebook

LinkedIn

Github

**Videos**

**Product Demo**

Paste Youtube/Vimeo Product Demo URL

Show how your product or prototype works in 1 minute or less.

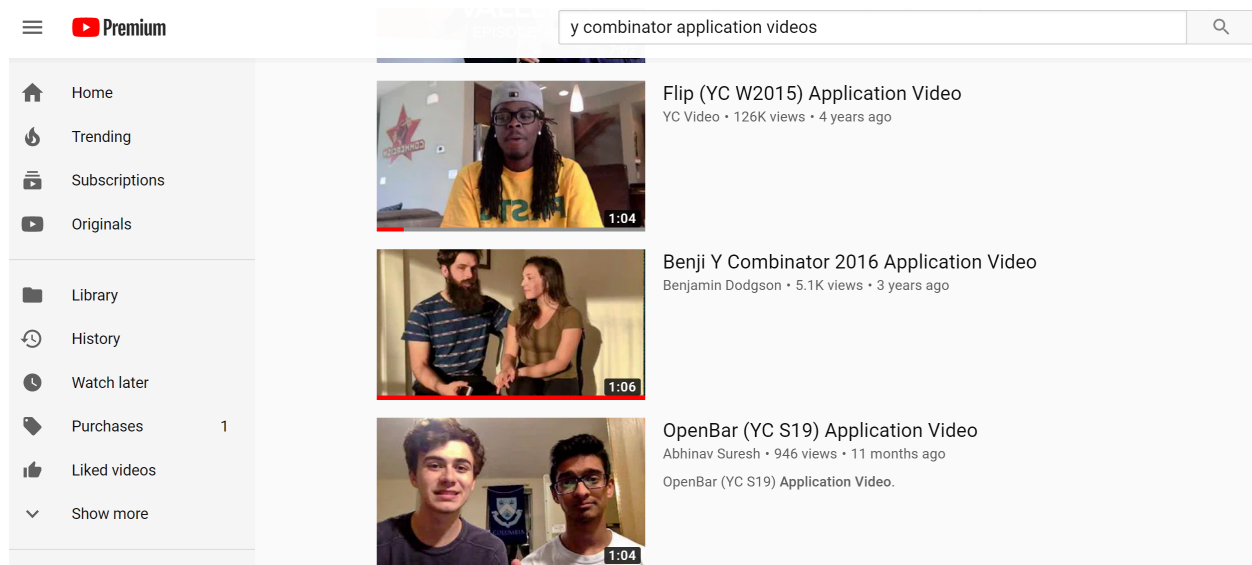
**Team \***

Paste Youtube/Vimeo Founders Intro URL

Introduce your team in 1 minute or less.

Use a public Youtube/Vimeo URL only (ex. [www.youtube.com/foo](#)). Do not password protect your video but non-public/unlisted is Ok.

**Figure IA.1. Examples of Accelerator Online Application Forms.** This figure shows screenshots of accelerators' online application forms. We show forms for Y Combinator, MassChallenge, 500 Startups, and Techstars. On each online application form, we highlight the question specifically asking for uploading pitch videos.



**Figure IA.2.** Screenshot of Search Results from YouTube

**Table IA.2**  
**Sample Description of Pitch Videos**

This table provides descriptive statistics on collected videos by accelerators that the applications are made to (Panel A) and by year (Panel B). We obtain pitch videos using an automatic searching script for two public video-sharing websites, YouTube and Vimeo. Integrated with query APIs, our web crawler returns a list of video indices according to a set of predefined keywords, which include but are not limit to the names of these accelerators, “startup accelerator application video”, “accelerator application videos” and so on. We first obtain the full list of potential videos returned by each keyword search (there is a limitation of returned videos by YouTube), and then filter the potential videos by a combination of different conditions on video info obtained along with the video itself. Filtering variables include but are not limit to data format, duration, title, and annotation. We also obtain additional videos from accelerators’ websites. Panel A reports the number of videos submitted to each accelerator and the proportion of each accelerator in the full sample. Panel B reports the breakdown by application year (typically the year of video uploading).

Panel A: Breakdown by Accelerators and Investment Status

Accelerator	Videos #	Accelerator Invested	Website Active	In Crunchbase	In PitchBook
500 Startups	33	1	15	19	8
AngelPad	83	2	33	36	18
MassChallenge	166	56	129	113	79
Techstars	136	3	67	53	21
Y Combinator	713	35	363	238	91
YC Fellowship	8	0	2	3	0
Total	1,139	97	609	462	217
% of Full Sample	100%	8.52%	53.47%	40.56%	19.05%

Panel B: Breakdown by Years

Accelerator	<=2012	2013	2014	2015	2016	2017	2018	2019
500 Startups	1	1	7	7	2	8	5	2
AngelPad	11	7	13	4	12	14	21	1
MassChallenge	4	9	4	13	34	33	34	35
Techstars	9	17	12	15	8	30	32	13
Y Combinator	10	31	29	82	67	110	164	220
YC Fellowship	0	0	0	8	0	0	0	0
Total	35	65	65	129	123	195	256	271
% of Full Sample	3.07%	5.71%	5.71%	11.33%	10.80%	17.12%	22.48%	23.79%

## A.2. *Details on Gathering Founder and Startup Information*

Founder-level control variables are constructed based on the information of the presenter(s) instead of the people listed as co-founders in external databases. To achieve this goal, our data collection processes involve comparing presenters' self-reported names and facial images with the names and pictures on individual profiles, and only information about the presenters are used. Below we offer more details, which we hope can mitigate any concerns.

- We obtain presenter names from self-introductions in pitches, video description text, and YouTube account names. These presenter names, along with startup names, are then used as keywords for searching on LinkedIn, our main data source to gather individual information.
- Among 1,139 startups in our sample, we are able to find the presenters on LinkedIn for 693 (61%) of them. For these startups, we collect information on presenters' educational backgrounds and work experiences. We code such information in an array of categorical variables, including whether presenters have a master's or a PhD degree, whether they attended an elite university, whether they have prior entrepreneurship experience, and whether they ever held a senior position in prior employment.
- For startups for which we are unable to find presenters' LinkedIn profiles, we construct the same array of categorical variables and code variables as the "missing" category. For example, the categorical variable of whether presenters have a master degree has three categories: "Yes", "No", and "Missing". We then add dummy variables that correspond to each categorical variable to our regressions as controls for team background.

In Table [IA.3](#), we conduct the following robustness tests, using the specification in Table [IV](#). First, we focus on the subsample of startups whose presenters can be found on LinkedIn. The effect of the Pitch Factor on the probability of receiving an investment remains significant. And the effect is larger relative to the full-sample estimate. Second, we add a dummy variable  $I(Has\ LinkedIn)$  to the specification. The dummy variable takes the value of one if we are able to find LinkedIn profiles of presenters and zero otherwise. The coefficients of the Pitch Factor remain stable. Meanwhile, the positive and significant coefficients of  $I(Has\ LinkedIn)$  indicate that teams whose presenters have LinkedIn profiles have a higher probability of receiving an investment.

**Table IA.3**  
**Investment Decisions and Missing LinkedIn Profiles**

Logit regressions, marginal effect. The analysis is obtained using the following model:

$$I(Invested) = \alpha + \beta \cdot X + \delta_{FE} + \varepsilon.$$

$I(Invested)$  takes a value of one if the startup team was chosen by the accelerator and zero otherwise. All pitch feature variables are standardized into a zero-mean variable with a standard deviation of one. All variables are identical to those in Table IV. The dummy variable  $I(Has\ LinkedIn)$  takes the value of one if we are able to find LinkedIn profiles of presenters and zero otherwise. Control variables include founders' education background (whether they have a master's or a PhD degree; whether they attended an elite university, defined as the U.S. News & World Report's Top 10), founders' prior work experience (whether they have prior entrepreneurship experience; whether they ever held a senior position in prior employment), team size, and video resolution. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent Var: $I(Invested)$					
Pitch-Factor	0.030*** (0.007)	0.026*** (0.007)	0.045*** (0.012)	0.040*** (0.012)	0.029*** (0.007)	0.026*** (0.007)
$I(Has\ LinkedIn)$					0.094*** (0.029)	0.113*** (0.031)
Observations	1,139	1,139	693	693	1,139	1,139
Pseudo $R^2$	0.193	0.302	0.158	0.251	0.229	0.302
Sample	Full	Full	Has LinkedIn	Has LinkedIn	Full	Full
Startup/Team Controls		Y		Y		Y
Content Controls		Y		Y		Y
Accelerator FE	Y	Y	Y	Y	Y	Y

## B. Method Appendix

This appendix provides more details on the steps to perform video analysis used in our paper. Compared to the more theoretical descriptions provided in Section II of the paper, this appendix proceeds with a more practical approach with information on our code structure, key functions, and notes on important steps.

### B.1. Examples of Visual, Vocal, and Verbal Features

Panel A: Example of High-Positivity Visual Features



Panel B: Example of Low-Positivity Visual Features



**Figure IA.3. Examples of Positive and Negative Visual Features.** This figure presents examples of a frame showing positive facial expressions (Panel A) and less-positive facial expressions (Panel B).



### Panel A: Example of High-Ability Pitch Script


Hi, I'm Vitali CEO of **Fitness** Lab. There are a lot of **fitness** apps in the world, but they all have the same problems. They offer their users random workout and diet plans. They use just **do** it marketing to push sales and they have low retention because people leave after three months. So we decided to fix it and **make** the **best fitness** app in the world and **power** it with artificial **intelligence**. We're going to offer our users highly personalized workout and diet plans by now. We have **launched** MVP. We have 50,000 installations and five thousand **active** users. We have received several **design award** including the one for the **best** interaction **design** by American Institute of graphic arts. So we have awesome **design** and very **efficient** technology. This is going to be more **efficient** than any living coach. We know that Y Combinator has a lot of connections with artificial **intelligence** businesses. That's why we're looking forward to your support. Thank you.

### Panel B: Example of High-Warmth Pitch Script


Hello, I'm Marcus and I'm Rebecca and **together** we're the proud founders of Fine Print Fighters LLC. We **help** expose small and misleading content in contracts. We **help** consumers make much more informed decisions during the purchasing process both pre and post purchase. We **like** to **help** the consumers gain **back** control of the purchasing process, and we **like** to create value well through our **pleasing** personalities as you can tell. Well, we look forward to working with angel pad, and we **appreciate** the opportunity in advance. Look forward to working with the staff and the rest of the constituents and hopefully be a good representation of what angel angel pad represents. So we **thank** you again in advance, and we look forward to speaking with you all and seeing you all soon. **Thank** you.

**Figure IA.4. Examples of High-Ability and Warmth Script.** This figure presents examples of startup pitch scripts with high-ability (Panel A) and high-warmth (Panel B) verbal features. The key ability words are highlighted in Panel A, and key warmth words are highlighted in Panel B.

Panel A: Visualized Example of High-Arousal Pitch

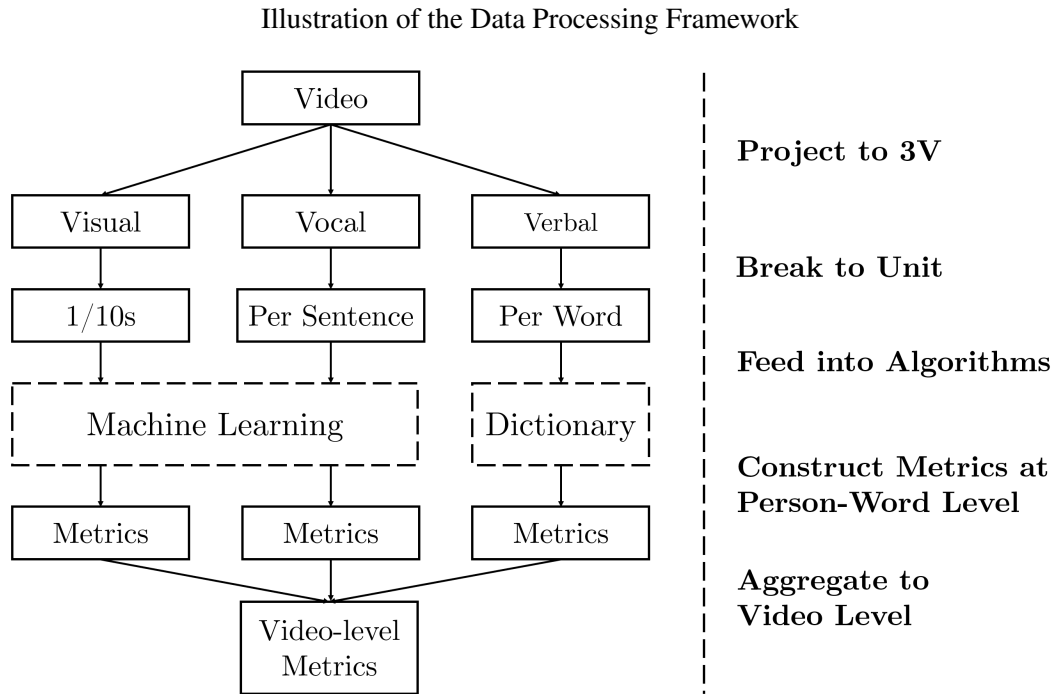


Panel B: Visualized Example of Low-Arousal Pitch



**Figure IA.5. Visualized Examples of High- and Low- Arousal Vocal Features.** This figure presents visualized examples of startup pitches with high-arousal (Panel A) and low-arousal (Panel B) vocal features. The visualization uses the waveform amplitude of those pitches. The high-arousal pitch can be downloaded from <https://www.dropbox.com/s/ipluo2w9tsszu2m/High%20Arousal%20Example.wav?dl=0>, and the low-arousal pitch can be downloaded from <https://www.dropbox.com/s/7igoqkjla72usdc/Low%20Arousal%20Example.wav?dl=0>.

## B.2. Video Processing Example



# Hu and Ma (2024) Video Processing Package

This example shows how to use `interactionvideo` package to process a video to study interpersonal interactions and communications.

Please refer to and cite the research paper: Hu and Ma (2024), "Persuading Investors: A Video-Based Study", available at [SSRN](#), by [Allen Hu](#) (UBC Sauder, allen.hu@sauder.ubc.ca) and [Song Ma](#) (Yale University and NBER, song.ma@yale.edu)

**For academic research purposes only.**

## Overview

The video processing involves the following steps:

1. Set up folders and check dependencies and requirements
2. Extract images and audios from a video using `pliers`
3. Extract text from audios using [Google Speech2Text API](#)
4. Process images (faces) using [Face++ API](#)
5. Process text using [Loughran and McDonald \(2011\)](#) Finance Dictionary and [Nicolas, Bai, and Fiske \(2019\)](#) Social Psychology Dictionary
6. Process audios using pre-trained ML models in `pyAudioAnalysis` and `speechemotionrecognition`
7. Aggregate information from 3V (visual, vocal, and verbal) to video level

## File Structure

- `example.py` and `example.ipynb` : step-by-step tutorials
  - We recommend you start from `example.ipynb`
- `interactionvideo` : main package folder
- `data` : data input folder
- `output` : output result folder
- `mlmodel` : pre-trained ML model folder
- `PythonSDK` : Face++ Python SDK folder, downloaded directly from [Github](#)

```
├── interactionvideo
│   ├── __pycache__
│   ├── prepare.py
│   ├── decompose.py
│   ├── faceppml.py
│   ├── googlem1.py
│   ├── textualanalysis.py
│   ├── audiom1.py
│   └── aggregate.py
```

```

├── utils.py
├── data
│   ├── example_video.mp4
│   └── VideoDictionary.csv
├── mlmodel
│   ├── pyAudioAnalysis
│   └── speechemotionrecognition
├── output
│   ├── audio_temp
│   ├── image_temp
│   └── result_temp
├── PythonSDK
├── README.md
├── requirement.txt
├── environment.yml
├── example.py
└── example.ipynb

```

## 1. Set up folders and check dependencies

```

In [3]: import os
        from os.path import join

        # Set the path
        RootPath = ''

        # Set your video file path here
        VideoFilePath = join(RootPath, 'data', 'example_video.mp4')

        # Set your work path here
        # Work path is where to store meta files and output files
        WorkPath = join(RootPath, 'output')

        # Set your path to the Google Cloud credential file here
        os.environ['GOOGLE_APPLICATION_CREDENTIALS'] = ''

        # Append path to ffmpeg
        FFMPEGPath = ''
        os.environ['PATH'] += os.pathsep + FFMPEGPath

```

```

In [4]: # Set up the folders
        from interactionvideo.prepare import setup_folder
        setup_folder(WorkPath)

        # check the requirements for interactionvideo
        from interactionvideo.prepare import check_requirements
        check_requirements()

```

decompose.py requirements satisfied.

faceppml.py requirements satisfied.

googleml.py requirements satisfied.

audioml.py requirements satisfied.

Out[4]: True

## 2. Extract images and audios from video

In [5]: `from interactionvideo.decompose import convert_video_to_images`

```
# Decompose the video into a stream of images  
# The default sampling rate is 10 frames per second  
# Find the output at WorkPath\image_temp  
convert_video_to_images(VideoFilePath, WorkPath)
```

Video is 70.12 seconds long.

100%|██████████| 702/702 [05:55<00:00, 1.98it/s]

Video is sampled to 702 images.

Video to images finished.

Out[5]: True

In [6]: `from interactionvideo.decompose import convert_video_to_audios`

```
# Decompose the video into audios  
# Find the output at WorkPath\audio_temp  
convert_video_to_audios(VideoFilePath, WorkPath)
```

[MoviePy] Writing audio in E:\Research Dropbox\Research\VideoProcessing\Code\Git Repo\202408\Python\output\audio\_temp\audio\_full.wav

100%|██████████| 1683/1683 [00:00<00:00, 2770.97it/s]

[MoviePy] Done.

Video to audios finished.

Out[6]: True

## 4. Process images (faces) using Face++ API

Use Face++ ML API to process images.

Get the key and secret from <https://www.faceplusplus.com>. If you register at <https://console.faceplusplus.com/register>, use <https://api-us.faceplusplus.com> as the server. If you register at <https://console.faceplusplus.com.cn/register>, use <https://api-cn.faceplusplus.com> as the server. The Python SDK of Face++ is included in this repo (Python SDK). Check the original package at <https://github.com/FacePlusPlus/facepp-python-sdk>.

In [7]: `from interactionvideo.faceppml import process_image_by_facepp`

```
# Use Face++ ML API to process images  
# Return csv files of facial emotion, gender, predicted age  
# Find the output  
# - WorkPath\result_temp\face_panel_facepp.csv (full returns from Face++)  
# - WorkPath\result_temp\face_panel.csv (clean results)
```

```
# Set your key, secret, and server here
FaceppKey = ''
FaceppSecret = ''
FaceppServer = 'https://api-us.faceplusplus.com'

facepp_result_df, facepp_result_clean_df = process_image_by_facepp(VideoFilePath, WorkPath,
                                                                    FaceppKey, FaceppSecret, FaceppServer)
```

Face++ API begins. 702 images to process.

100%|██████████| 702/702 [1:09:10<00:00, 5.91s/it]

Face++ API ends. 702 images processed.

In [8]: *# Check full returns from Face++*  
facepp\_result\_df.head(10)

Out[8]:

	ImageName	Onset	Offset	Duration	face_rectangle#top	face_rectangle#left	face_rectangle#width	face_
0	frame[0]	0.0	0.1	0.1	405	868	249	
1	frame[3]	0.1	0.2	0.1	406	867	250	
2	frame[6]	0.2	0.3	0.1	404	866	252	
3	frame[9]	0.3	0.4	0.1	403	867	253	
4	frame[12]	0.4	0.5	0.1	401	866	258	
5	frame[15]	0.5	0.6	0.1	405	867	261	
6	frame[18]	0.6	0.7	0.1	407	867	261	
7	frame[21]	0.7	0.8	0.1	404	869	258	
8	frame[24]	0.8	0.9	0.1	403	867	262	
9	frame[27]	0.9	1.0	0.1	402	868	262	

10 rows × 193 columns

In [9]: *# Check clean results*  
facepp\_result\_clean\_df.head(10)

Out[9]:

	Onset	Offset	Duration	Number of Faces	Gender	Age	Visual-Positive	Visual-Negative	Visual-Beauty
0	0.0	0.1	0.1	1	Male	31	0.00007	0.26985	0.430860
1	0.1	0.2	0.1	1	Male	33	0.00008	0.22857	0.406690
2	0.2	0.3	0.1	1	Male	30	0.00115	0.32879	0.413970
3	0.3	0.4	0.1	1	Male	28	0.00152	0.33477	0.402910
4	0.4	0.5	0.1	1	Male	28	0.00040	0.92615	0.415210
5	0.5	0.6	0.1	1	Male	26	0.00734	0.98612	0.447690
6	0.6	0.7	0.1	1	Male	30	0.00187	0.79954	0.449705
7	0.7	0.8	0.1	1	Male	32	0.00021	0.09503	0.449665
8	0.8	0.9	0.1	1	Male	29	0.00095	0.61223	0.451375
9	0.9	1.0	0.1	1	Male	29	0.00046	0.05613	0.468775

### 3. Extract text from audios using Google Speech2Text API

Set up your Google Cloud environment following

- <https://cloud.google.com/python>
- <https://cloud.google.com/storage/docs/quickstart-console>
- <https://developers.google.com/workspace/guides/create-credentials>
- <https://cloud.google.com/storage/docs/creating-buckets>

Create a Google Cloud Storage bucket.

```
In [12]: from interactionvideo.googleml import upload_audio_to_googlecloud
```

```
# Set your Google Cloud Storage bucket name here
GoogleBucketName = ''

# Upload audio file to Google Cloud Storage
upload_audio_to_googlecloud(WorkPath, GoogleBucketName)
```

Uploaded the audio file to Google Cloud.

```
Out[12]: True
```

```
In [13]: from interactionvideo.googleml import convert_audio_to_text_by_google
```

```
# Use Google Speech2Text API to convert audio to text
# Return a txt file of full speech script and a csv file of text and punctuation
# Find the output at
# - WorkPath\result_temp\script_google.txt (full speech script)
# - WorkPath\result_temp\text_panel_google.csv (text panel from Google)
google_result_text, google_result_df = convert_audio_to_text_by_google(WorkPath, GoogleBuck
```



Google Speech2Text begins. 70.12 seconds audio to process.

Google Speech2Text ends. 70.12 seconds audio processed.

```
In [14]: # Check full speech script from Google
print(google_result_text)
```

Hello, everyone. First of all, we will like to thank you for your interest, in our research. In this paper, we try to understand how human interaction features such as facial expressions, vocal emotions, and word choices, might influence economic agents decision-making. In order to study this question. Empirically, we build an empirical approach that uses videos of human interactions as data input, and, and machine learning based algorithms as the tool. We applied a couple approached in a setting where early stage Startup pitch, Venture capitalists, for early stage funding. We find that pitch features along visual vocal and verbal Dimensions, all matter for the probability of receiving funding. And we also show that this event impact is largely due to interaction, induced biases, rather than that interactions provide additional valuable information. The empirical structure that you see, in this code example, can hopefully help you to get started with using video in other important settings. Such as Interviews. Classroom recordings. Among many other exciting things will look forward to hearing your feedback and reading about your research. Thank you.

```
In [15]: # Check text panel from Google
google_result_df.head(10)
```

```
Out[15]:
```

	Text	Onset	Offset	Duration	Sentence End
0	Hello,	0.4	0.6	0.2	True
1	everyone.	0.6	1.2	0.6	True
2	First	1.2	1.5	0.3	False
3	of	1.5	1.5	0.0	False
4	all,	1.5	1.9	0.4	True
5	we	1.9	2.0	0.1	False
6	will	2.0	2.2	0.2	False
7	like	2.2	2.3	0.1	False
8	to	2.3	2.4	0.1	False
9	thank	2.4	2.7	0.3	False

## 5. Process text using LM and NBF Dictionaries

Use Loughran-McDonald (2011) Finance Dictionary (LM) to construct verbal positive and negative.

For more details, please check <https://sraf.nd.edu/textual-analysis/resources>.

Use Nicolas, Bai, and Fiske (2019) Social Psychology Dictionary (NBF) to construct verbal warmth and ability.

For more details, please check <https://psyarxiv.com/afm8k>.

```
In [16]: from interactionvideo.textualanalysis import process_text_by_dict

# Set LM-NBF dictionary path
DictionaryPath = join(RootPath, 'data', 'VideoDictionary.csv')

# Dictionary-based textual analysis to get verbal measures
# (e.g., verbal positive, negative, warmth, ability)
# Return csv files of verbal positive, negative, warmth, and ability
# Find the output at
# - WorkPath\result_temp\text_panel.csv
text_result_df = process_text_by_dict(WorkPath, DictionaryPath)
```

LM and NBF Dictionaries loaded.

Dictionary-based textual analysis begins. 181 words to process.

Dictionary-based textual analysis ends. 181 words processed.

```
In [17]: # Check text panel from Dictionary
text_result_df.head(10)
```

```
Out[17]:
```

	Text	Onset	Offset	Duration	Sentence End	Verbal- Negative	Verbal- Positive	Verbal- Warmth	Verbal- Ability
0	Hello,	0.4	0.6	0.2	True	0.0	0.0	0.0	0.0
1	everyone.	0.6	1.2	0.6	True	0.0	0.0	0.0	0.0
2	First	1.2	1.5	0.3	False	0.0	0.0	0.0	0.0
3	of	1.5	1.5	0.0	False	0.0	0.0	0.0	0.0
4	all,	1.5	1.9	0.4	True	0.0	0.0	0.0	0.0
5	we	1.9	2.0	0.1	False	0.0	0.0	0.0	0.0
6	will	2.0	2.2	0.2	False	0.0	0.0	0.0	0.0
7	like	2.2	2.3	0.1	False	0.0	0.0	1.0	0.0
8	to	2.3	2.4	0.1	False	0.0	0.0	0.0	0.0
9	thank	2.4	2.7	0.3	False	0.0	0.0	1.0	0.0

## 6. Process audios by pre-trained ML models

Construct vocal arousal and vocal valence from pre-trained SVM ML models in `pyAudioAnalysis`.

The pre-trained models are located at `mlmodel\pyAudioAnalysis`

- `svmSpeechEmotion_arousal`
- `svmSpeechEmotion_arousalMEANS`
- `svmSpeechEmotion_valence`
- `svmSpeechEmotion_valenceMEANS`

For more details, please check <https://github.com/tyiannak/pyAudioAnalysis/wiki/4.-Classification-and-Regression>.

Construct vocal positive and vocal negative from pre-trained LSTM ML models in `speechemotionrecognition`.

The pre-trained models are located at `mlmodel\speechemotionrecognition`

- `best_model_LSTM_39.h5`

For more details, please check <https://github.com/harry-7/speech-emotion-recognition>.

Note: `speechemotionrecognition` requires tensorflow and Keras.

```
In [18]: from interactionvideo.audioml import process_audio_by_pyAudioAnalysis

# Set the model path
pyAudioAnalysisModelPath = join(RootPath, 'mlmodel', 'pyAudioAnalysis')

# Construct vocal arousal and vocal valence
# Find the output at
# - WorkPath\result_temp\audio_panel_pyAudioAnalysis.csv
audio_result_df1 = process_audio_by_pyAudioAnalysis(WorkPath, pyAudioAnalysisModelPath)
```

Using TensorFlow backend.

pyAudioAnalysis vocal emotion analysis begins. 70.12 seconds audio to process.

pyAudioAnalysis ML model loaded.

pyAudioAnalysis vocal emotion analysis ends. 70.12 seconds audio processed.

```
In [19]: # Check audio panel from pyAudioAnalysis
audio_result_df1.head()
```

```
Out[19]:
```

	Onset	Offset	Duration	Vocal-Arousal	Vocal-Valence
0	0	70.12	70.12	0.404801	-0.015317

```
In [20]: from interactionvideo.audioml import process_audio_by_speechemotionrecognition

# Set the model path
speechemotionrecognitionModelPath = join(RootPath, 'mlmodel', 'speechemotionrecognition')

# Construct vocal positive and vocal negative
# Find the output at
# - WorkPath\result_temp\audio_panel_speechemotionrecognition.csv
audio_result_df2 = process_audio_by_speechemotionrecognition(WorkPath, speechemotionrecognitionModelPath)
```

speechemotionrecognition vocal emotion analysis begins. 70.12 seconds audio to process.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:66: The name tf.get\_default\_graph is deprecated. Please use tf.compat.v1.get\_default\_graph instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:541: The name tf.placeholder is deprecated. Please use tf.compat.v1.placeholder instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:4432: The name tf.random\_uniform is deprecated. Please use tf.random.uniform instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:148: The name tf.placeholder\_with\_default is deprecated. Please use tf.compat.v1.placeholder\_with\_default instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:3733: calling dropout (from tensorflow.python.ops.nn\_ops) with keep\_prob is deprecated and will be removed in a future version.

Instructions for updating:

Please use `rate` instead of `keep\_prob`. Rate should be set to `rate = 1 - keep\_prob`.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\optimizers.py:793: The name tf.train.Optimizer is deprecated. Please use tf.compat.v1.train.Optimizer instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:3657: The name tf.log is deprecated. Please use tf.math.log instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\tensorflow\_core\python\ops\nn\_impl.py:183: where (from tensorflow.python.ops.array\_ops) is deprecated and will be removed in a future version.

Instructions for updating:

Use tf.where in 2.0, which has the same broadcast rule as np.where

Model: "sequential\_1"

Layer (type)	Output Shape	Param #
=====		
lstm_1 (LSTM)	(None, 128)	86016
-----		
dropout_1 (Dropout)	(None, 128)	0
-----		
dense_1 (Dense)	(None, 32)	4128
-----		
dense_2 (Dense)	(None, 16)	528
-----		
dense_3 (Dense)	(None, 4)	68
=====		
Total params: 90,740		
Trainable params: 90,740		
Non-trainable params: 0		

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:190: The name tf.get\_default\_session is deprecated. Please use tf.compat.v1.get\_default\_session instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:197: The name tf.ConfigProto is deprecated. Please use tf.compat.v1.ConfigProto instead.

mpat.v1.ConfigProto instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:203: The name tf.Session is deprecated. Please use tf.compat.v1.Session instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:207: The name tf.global\_variables is deprecated. Please use tf.compat.v1.global\_variables instead.

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:216: The name tf.is\_variable\_initialized is deprecated. Please use tf.compat.v1.is\_variable\_initialized instead.

None

WARNING:tensorflow:From d:\Software\Anaconda3\envs\interactionvideo\lib\site-packages\keras\backend\tensorflow\_backend.py:223: The name tf.variables\_initializer is deprecated. Please use tf.compat.v1.variables\_initializer instead.

speechemotionrecognition ML model loaded.

speechemotionrecognition vocal emotion analysis ends. 70.12 seconds audio processed.

```
In [21]: # Check audio panel from speechemotionrecognition
audio_result_df2.head()
```

```
Out[21]:
```

	Onset	Offset	Duration	Vocal-Positive	Vocal-Negative
0	0	70.12	70.12	0.487744	0.0068

## 7. Aggregate information from 3V to video level

```
In [22]: from interactionvideo.aggregate import aggregate_3v_to_video
```

```
# Aggregate 3V information
# Find the output at
# - WorkPath\result_temp\video_panel.csv
video_result_df = aggregate_3v_to_video(WorkPath)
```

3V to video aggregation finished.

```
In [23]: # Check video panel
video_result_df.T
```

Out[23]:

0	
Number of Faces	1.0
Gender	Male
Age	32.0
Visual-Positive	0.014174
Visual-Negative	0.44359
Visual-Beauty	0.450598
Vocal-Positive	0.49
Vocal-Negative	0.01
Vocal-Arousal	0.4
Vocal-Valence	-0.02
Verbal-Positive	0.01105
Verbal-Negative	0.005525
Verbal-Warmth	0.033149
Verbal-Ability	0.038674

### B.3. Textual Analysis on Pitch Content

In this appendix, we provide more technical details on the construction of informational content measures for the pitches.

**Measures of idea novelty based on textual similarity.** We measure the novelty of ideas in video pitches by comparing their textual content with business descriptions of startups and public firms extant around the same time. The idea is that if the pitch of the focal startup is different from existing businesses (i.e., not a me-too startup) but could be influential in the future (i.e., the idea will have some traction), we consider the pitching startup to be more novel. [Kelly et al. \(2021\)](#) at the *AER: Insights* takes a similar empirical strategy to measure the technological novelty of patents.

To implement this idea, we obtain a panel of business descriptions of existing startups from PitchBook and of publicly-traded firms from the business description section (Item 1) from 10-K filings of these firms. Combining these data, we observe the business descriptions of startups founded each year and the descriptions of public firms each year.

Our measure construction process closely follows that of [Kelly et al. \(2021\)](#). For a focal startup  $i$  in our pitch sample that applied to an accelerator in year  $t$ , we construct its idea novelty measure in three steps.<sup>1</sup>

- **Step 1:** We calculate “backward textual similarity” as the average textual similarity (more on this below) between  $i$ ’s pitch script and business descriptions of all startups that were first financed by early-stage VCs before or in year  $t$ . A low backward textual similarity indicates that startup  $i$ ’s idea is distinct from the business models of previously and contemporaneously funded startups.
- **Step 2:** We calculate “forward textual similarity” as the average textual similarity between  $i$ ’s pitch script and business descriptions of all startups that were first financed by early-stage VCs after year  $t$ . A high forward textual similarity indicates that startup  $i$ ’s idea is similar to the business models of startups funded in the future.
- **Step 3:** We calculate the novelty measure using both the backward and forward textual similarities—dividing the forward one by the backward one. Together, a high forward-to-backward ratio indicates a high novelty for startup  $i$ ’s idea: it is different enough from previous

---

<sup>1</sup>To keep the description concise, we skipped the standard processes of textual cleaning in this description. We are happy to provide more details if needed.

ideas but is potentially impactful for the future. The same logic applies to the measure using public firms’ business descriptions as benchmarks.

A key component in the calculation above is the definition of textual similarity. We calculate textual similarities using both BERT and Bag-of-Words (BoW), and the results are robust to both.

- To quantify the information embedded in text, we first need to represent the textual data in a numerical format. We use Bidirectional Encoder Representations from Transformers (BERT), a state-of-the-art NLP model of word embeddings that maps words into vectors of real numbers. BERT proves to be superior in many NLP tasks (Devlin et al., 2018) and has been increasingly used in economic studies (Gorodnichenko et al., 2023). We use “all-mpnet-base,” the current best-performing version of BERT in sentence embedding<sup>2</sup>.
- As a robustness test, we use the “bag of words” (BoW) representation with the “term-frequency-inverse-document-frequency” (TF-IDF) weighting scheme and obtain similar results. For each video pitch or business description, we use BERT to transform it to a vector. We then define the textual similarity between video pitch and business description as the cosine similarity between each pair of vectors.

**Dictionary-based measures of pitch content.** We use a dictionary-based approach to directly capture the topics that are discussed in video pitches. We focus on the topics that are most relevant in the setting of early-stage startup financing. These topic categories include concrete numbers, cash flow, competition, employment, readiness, technology, data, and AI. We compile a list of keywords that are representative of these topic categories. For example, the keywords for the “cash flow” category include “sale(s)”, “revenue(s)”, and “profit(s)”, among others, which capture whether a startup discusses the profitability in the video pitch. The “technology” category is concerned with whether the pitch explicitly discusses the technologies or patents. We define the category of concrete numbers as all numbers mentioned in video pitches. Table IA.4 shows a complete list of categories and keywords.

For each video pitch, we examine whether the keywords of each topic are included. The dummy variable of each topic takes a value of one if any keyword of that topic appears in the content of

---

<sup>2</sup>For a complete list of BERT versions, see [https://www.sbert.net/docs/pretrained\\_models.html](https://www.sbert.net/docs/pretrained_models.html)



a pitch. For example, the measure “Competition” has a mean of 0.06, which indicates that 6% of startups in our sample discuss competition explicitly in their video pitches.

**LIWC.** We use LIWC to extend the word categories of our dictionary-based approach. LIWC is widely used in computational linguistics and includes word categories that capture soft information and psychological meanings of text ([Tausczik and Pennebaker, 2010](#)). Over 20,000 scientific articles have already been published using LIWC. Similar to our practice above, for each LIWC category, we calculate its percentage of total words within a video pitch.

To complement the word categories in the finance dictionary ([Loughran and McDonald, 2011](#)), the social psychology dictionary ([Nicolas et al., 2019](#)), and our startup financing word list described above, we focus on communication styles (e.g., concrete and informal words) and time orientations (e.g., past, present, or future focus) in LIWC.

**Table IA.4**  
**List of Keywords of Content Control Categories**

This table lists the keywords for constructing dictionary-based measures of informational content controls.

Category	Keywords	Category	Keywords	Category	Keywords
Cash Flow	sale(s)	Technology	patent(s)	Data and AI	digitalization
	revenue(s)		innovation(s)		digitalize(s)
	profit(s)		invention(s)		digitally
	profitability		inventor(s)		digitize(s)
	income(s)		technique(s)		digitized
	earning(s)		technology(ies)		digitizing
	cash flow(s)		technological		program
Employment	employ(s)	Competition	compete(s)		programmed
	employing		competing		programming
	employed		competition(s)		programmer(s)
	employment		competitive		programmatic
	employee(s)		competitiveness		programmable
	employer(s)		competitor(s)		artificial intelligence
	recruit(s)	Data and AI	data		machine learning
	recruited		database		
	recruiting		information		
	recruiter(s)		analysis		
Readiness	recruitment		analyses		
	prototype(s)		analytic		
	prototyping		analytical		
	customer(s)		analytics		
	commercialize(s)		analyze(s)		
	commercialized		analyzed		
	commercialise(s)		analyzing		
	commercialised		developer(s)		
	commercialization		digital		

## C. Analysis Using the University Sample

To examine the robustness of our results, we analyze a new sample from the Yale Tsai Center for Innovative Thinking (CITY) using the same video analysis technique and empirical strategy. We show below that the empirical results of the Yale Sample Analysis are very similar to our original results in the paper, both in terms of economic magnitudes and statistical significance. With this administrative-level data set, we also perform a test on the sample selection problem which arises from the initial decision to submit and publicize a video pitch. We find that the availability of a video is neither related to measures of pitch features nor investment decisions, which suggests that the selection issue may not be a major concern.

### C.1. Data

Yale Tsai CITY is an institute at Yale University that aims to inspire innovation and entrepreneurship. Yale Tsai CITY runs an accelerator program, taking applications three times a year (Spring, Summer, Fall) from startup teams formed among Yale students. The application process is very similar to the accelerator programs studied in our main analysis (albeit on a much smaller scale). The applications are reviewed, and accepted teams receive an investment of \$2,000 and additional resources such as mentorship, expert services (on accounting, legal services, and communication), and community activities—again, similar to the commercial accelerators.

Our data include all 316 Yale Tsai CITY accelerator applications between Fall 2018 and Fall 2020. For each application, we obtain the information submitted through the online form and the pitch video submitted together with the application. Among 316 startups in our sample, 166 (53%) include videos in their applications and 150 (47%) do not. Among 166 startups with videos included in their applications, 61 (37%) are funded by Yale Tsai CITY. For those 150 startups that do not include a video in their applications, 28 (19%) are funded by Yale Tsai CITY. Together, 89 (28%) are funded in the full sample.

**[Insert Table [IA.5](#) Here.]**

Table [IA.5](#) presents summary statistics of startups and videos, and it shows that the Yale sample is quite similar to our main sample in the paper. Similar to our original sample, the majority of these startups are still in an early stage—about 90% of them have not launched their products yet at the

time of application. The videos in the Yale Tsai CITY sample are slightly longer (111 seconds on average) than those in our original sample (83 seconds on average) since Yale Tsai CITY does not have a hard restriction on the length of video pitches, while some accelerators in our original sample require videos to be less than one minute. In general, these two samples are quite comparable in terms of startup characteristics and video features.

### *C.2. Robustness of Main Results*

Table [IA.6](#) repeats the video analysis procedure used on our original sample. Specifically, we then estimate the same model in Table [IV](#) to examine the relationship between Pitch Factor and the investment decisions by the accelerator. We show that in this Yale sample, Pitch Factor is significantly positively correlated with the probability for a startup to obtain funding from the accelerator. Our results are robust to alternative specifications with different sets of controls, such as team and video controls and time and startup-stage fixed effects. In terms of economic magnitude, take the coefficient in column (4) for example—a one-standard-deviation higher in Pitch Factor is associated with a 9.8 percentage point higher funding probability, which is equivalent to a 34.8 percent increase from the baseline funding rate of 28.16 percentage points. Such an estimate of 34.8 percent is very close to the one estimated in our original sample, which is 35.2 percent.

**[Insert Table [IA.6](#) Here.]**

### *C.3. Sample Selection Analysis*

One potential sample selection problem could arise from the initial decision to submit and publicize a video pitch. Since we collect the main sample from public domains and are unable to observe videos that were uploaded to private domains, our original sample is unable to speak to such a concern of sample selection. To better address this issue, we use the Yale Tsai CITY sample, which includes links to videos that are on private domains as well. Such a sample allows us to test whether the decision of making videos public is related to the key variables in our analysis. As shown in Table [IA.7](#), making the video public is unrelated to Pitch Factor, whether the startup is funded by the accelerator, and whether the startup receives funding from other investors and how much it receives.

**[Insert Table [IA.7](#) Here.]**

A key difference between the Yale sample and the main sample in the paper, however, is that student-led startup teams in the former have a lower probability of turning into more serious startups post-graduation, limiting our ability to study their long-term results in employment, obtaining VC funding, etc.

**Table IA.5**  
**Summary Statistics of Videos and Startups: Yale Tsai CITY**

This table provides descriptive statistics of pitch videos and the underlying startups in our sample of Yale Tsai CITY. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles. Panel A reports the total number of teams in each stage and in each period. Panel B reports characteristics of videos and teams.

Panel A: Breakdown by Periods and Stages

Period	Well-developed			Alpha/		Total	% of Full Sample
	Ideation	Idea	Prototyping	Beta	Launched		
Fall 2018	14	9	11	6	2	42	13.29%
Spring 2019	11	12	25	13	3	64	20.25%
Summer 2019	1	8	24	11	5	49	15.51%
Fall 2019	10	8	14	2	4	38	12.03%
Spring 2020	3	7	17	4	3	34	10.76%
Summer 2020	7	1	21	15	6	50	15.82%
Fall 2020	10	0	7	13	9	39	12.34%
Total	56	45	119	64	32	316	100%
% of Full Sample	17.72%	14.24%	37.66%	20.25%	10.13%	100%	

Panel B: Summary Statistics of Video and Startups

	N	Mean	STD	25%	50%	75%
Duration (second)	166	111.72	27.30	100.20	118.03	121.32
Video Size (MB)	166	24.38	44.52	7.16	15.66	21.35
Number of Words	166	303.83	73.52	259.00	304.50	353.00
Number of Sentences	166	33.63	9.10	28.00	33.00	38.00
Number of People	316	2.14	1.26	1.00	2.00	3.00
I( <i>Invested</i> )	316	0.28	0.45	0.00	0.00	1.00
I( <i>Other Funding</i> )	316	0.18	0.39	0.00	0.00	0.00
Other Funding Amount (\$000)	316	9.34	58.28	0.00	0.00	0.00

**Table IA.6**  
**Investment Decisions: Yale Tsai CITY**

Logit regressions, marginal effect. The analysis is obtained using the following model:

$$I(Invested) = \alpha + \beta \cdot X + \delta_{FE} + \varepsilon.$$

$I(Invested)$  takes a value of one if the startup team was chosen by the accelerator and zero otherwise. All pitch feature variables are standardized into a zero-mean variable with a standard deviation of one. Control variables include team size, video resolution, and content. Standard errors clustered at the period level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
	Dependent Var: $I(Invested)$				
Pitch Factor	0.091** (0.036)	0.091*** (0.035)	0.094** (0.036)	0.091*** (0.030)	0.103*** (0.024)
Observations	166	166	166	166	166
Pseudo $R^2$	0.027	0.028	0.125	0.159	0.290
Team/Video Controls	N	Y	Y	Y	Y
Content Controls	N	N	Y	Y	Y
Period FE	N	N	N	Y	Y
Stage FE	N	N	N	N	Y

**Table IA.7**  
**Sample Selection: Yale Tsai CITY**

Logit regressions, marginal effect. This table investigates the sample selection issue of the video sample. The analysis is restricted to the sample of startups that includes videos in their applications.  $I(\text{Public Videos})$  takes a value of one if the startup team uploads its video to a public domain and zero otherwise.  $I(\text{Invested})$  takes a value of one if the startup team was chosen by the accelerator and zero otherwise.  $I(\text{Other Funding})$  takes a value of one if the startup team has received funding from other investors at the time of application and zero otherwise. Other Funding Amount is the inverse hyperbolic sine of total amount of investment that a startup has raised from other investors at the time of application. All pitch feature variables are standardized into a zero-mean variable with a standard deviation of one. Standard errors clustered at the period level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
	<hr/> I(Public Video) <hr/>				
I(invested)	−0.088 (0.092)				−0.126 (0.112)
I(Other Funding)		−0.011 (0.049)			0.333 (0.256)
Other Funding Amount			−0.004 (0.004)		−0.039 (0.027)
Pitch Factor				0.002 (0.015)	0.016 (0.023)
Observations	166	166	166	166	166
Pseudo $R^2$	0.159	0.149	0.150	0.149	0.176
Content Controls	Y	Y	Y	Y	Y
Period FE	Y	Y	Y	Y	Y
Stage FE	Y	Y	Y	Y	Y



## D. Appendix: MTurk Rating Survey

This appendix presents details of our survey designs. The goal of these exercises is to bridge our ML-algorithm that rates pitch videos with the traditional approach of using human raters.


Both exercises take the form of an online survey that participants complete using their own electronic devices (e.g., computers and tablets), and they are distributed through Amazon Mechanical Turk (MTurk). In both surveys, we require the participants to be located in the U.S. and to be identified as masters at completing our types of tasks by the MTurk platform through its statistical performance monitoring. The experiments recruit 115 and 100 participants respectively. Our experiments on MTurk provide relatively high payments compared to the MTurk average to ensure quality responses.

Sample survey designs are attached toward the end of this appendix.

### *D.1. Survey 1: Rating on Pitch Positivity*

In the first survey, we elicit ratings of positivity from MTurkers. In each survey, a respondent is allocated a random set of six pitch videos. For each video, we first mandate the completion of watching the full video, and the respondent is not able to skip the video before answering the rating questions. Then, on the next survey screen, we elicit the rating of positivity, defined as passion, enthusiasms, based on the video just watched. The rating is on a 1-9 scale with nine choices. The evaluations of the videos are completed one by one, and ratings may not be revised after moving to the next video.

We then compare the ratings from MTurkers with the Pitch Factor. Figure [IA.6](#) shows the binned scatter plot of the relation between the two variables. The clearly positive correlation provides the first assurance of the validity of the ML-generated measure. In a regression analysis, as shown in Table [IA.8](#), we also show a strong correlation between the two variables.



**Figure IA.6.** Pitch Factor and Respondent-Rated Positivity

**Table IA.8**  
**Pitch Factor and Respondent-Rated Positivity**

	(1)	(2)
	Pitch Factor	
Respondent-Rated Positivity	0.062** (0.028)	0.088*** (0.034)
Observations	690	690
$R^2$	0.011	0.167
Respondent FE	No	Yes

## D.2. Survey 2: Comparing Pitches

In our second and separate survey, we ask MTurker respondents to compare pitch positivity in pairs of randomly-drawn videos. By asking respondents to directly compare pitches, we mitigate noise that could arise from the rating survey in Survey 1 due to the small sample—such as the impact of the order of videos and individual fixed effects in interpreting scales, among others.

In this survey, each respondent is allocated four pairs of videos. For each of these random pairs, we require both videos, clearly labeled as “Video 1” and “Video 2,” to be completely watched. Then on the next screen, the respondents are asked to choose the pitch video that gives them the

more positive impression (passionate, enthusiastic). Finally, we evaluate the consistency between our ML-based ranking and the human ranking. In other words, does the algorithm pick the same winners as the raters?

We find that the same winner is picked with nearly 89.5% consistency. Interestingly, we also find strong disagreement among MTurker raters themselves when the two videos in the same pair have close algorithm-generated Pitch Factors. In other words, our method seems to be able to provide a more decisive ranking when there are high levels of noise.

## Video Pitch Experiment Introduction

---

This survey will take you about 10 minutes. You will get a base payment of \$3 as long as you finish this survey. We will also award you bonus payment (up to \$3), which is determined by how well you did in the survey.

During the survey, you are going to watch 6 videos where company founders are describing their startup. You will then rate how positive (e.g. passionate, happy, enthusiastic) each video is on four dimensions: **facial expressions, voices, word choices, and overall**.

**Please get your audio device (e.g., earphone and computer speaker) ready now.**

***Note: The submission button will appear only after you watch the video. If the submission button does not appear even after you watch the video, please wait several seconds and do not reload the web page.***

---

## Video Pitch -Kru865yB-M (Example)

---

ConquerX (YC Winter 2019)



Please watch the video. You will then rate how positive (e.g. passionate, happy, enthusiastic) this video is on four dimensions: **facial expressions, voices, word choices, and overall**.  
(The submission button will appear after the video is played.)

---

## Video Pitch -Kru865yB-M Question (Example)

---

Which of the following industry or industries best describe the business of this startup?

- ☐ Consumption Goods
  - ☐ Health Care
  - ☐ Information Technology
  - ☐ Consumer Services
  - ☐ Industrials
- 

What is your rating for the **overall positivity** of this video?

	1	2	3	4	5	6	7	8	9	
Most negative	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	Most positive

---

What is your rating for the **visual positivity** of this video?

	1	2	3	4	5	6	7	8	9	
Most negative	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	Most positive

---

What is your rating for the **vocal positivity** of this video?

	1	2	3	4	5	6	7	8	9	
Most negative	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	Most positive

---

What is your rating for the **verbal positivity** of this video?

	1	2	3	4	5	6	7	8	9	
Most negative	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	Most positive

---

---

## Questions on Basic Information

---

What is your year of birth? (e.g., 1990)

---

Choose one or more races that you consider yourself to be:

- |  |   |
|--|---|
| <input type="checkbox"/> White                     | <input type="checkbox"/> Hispanic or Latino |
| <input type="checkbox"/> Asian                     | <input type="checkbox"/> Other              |
| <input type="checkbox"/> Black or African American |   |

---

What is your gender?

- ☐ Male
- ☐ Female
- ☐ Other

---

What is the highest level of school you have completed or the highest degree you have received?

- ☐ Less than High School
- ☐ High School
- ☐ College
- ☐ Graduate or Professional (JD, MD)
-

---

## Ending

---

This is the end of the survey. Thank you for your valuable time.

To obtain your payment, please input your unique ID below to MTurk.

Here is your unique ID: \${e://Field/Random%20ID}. Copy this value to paste into MTurk.

When you have copied this ID, please click the Submit button to submit your answers.

Powered by Qualtrics



## Video Pitch Experiment Introduction

---

This survey will take you about **15** minutes. You will get a base payment of **\$3** as long as you finish this survey. We will also award you bonus payment (up to **\$3**), which is determined by how well you did in the survey.

During the survey, you are going to watch **4 pairs** of videos where company founders are describing their startup. You will then select which video is more positive (e.g. passionate, happy, enthusiastic) on four dimensions: **facial expressions, voices, word choices, and overall**.

**Please get your audio device (e.g., earphone and computer speaker) ready now.**

***Note: The submission button will appear only after you watch both videos. If the submission button does not appear even after you watch the video, please wait several seconds and do not reload the web page.***

---

## Video Pitch le3s6qSV1Ck and n4d1TXm-RUk (Example)

---

### Video 1

AirOffice - 1 minute for Ycombinator



### Video 2

Green energy exchange video y combinator



Please watch both videos. You will then select which video is more positive (e.g. passionate, happy, enthusiastic) on four dimensions: **facial expressions, voices, word choices, and overall.**

(The submission button will appear after both videos are played.)

---

### Video Pitch le3s6qSV1Ck and n4d1TXm-RUk Question (Example)

---

Which of the following industry or industries best describe the business of the startup in video 1?

- ☐ Consumption Goods
  - ☐ Health Care
  - ☐ Information Technology
  - ☐ Consumer Services
  - ☐ Industrials
- 

Which of the following industry or industries best describe the business of the startup in video 2?

- ☐ Consumption Goods
  - ☐ Health Care
  - ☐ Information Technology
  - ☐ Consumer Services
  - ☐ Industrials
- 

Which video is more positive in terms of **overall positivity**?

	Video 1	Video 2
Overall Positivity	<input type="radio"/>	<input type="radio"/>

---

Which video is more positive in terms of **visual positivity**?

	Video 1	Video 2
Visual Positivity	<input type="radio"/>	<input type="radio"/>

---

Which video is more positive in terms of **vocal positivity**?

	Video 1	Video 2
Vocal Positivity	<input type="radio"/>	<input type="radio"/>

---

Which video is more positive in terms of **verbal positivity**?

	Video 1	Video 2
Verbal Positivity	<input type="radio"/>	<input type="radio"/>

---

## Questions on Basic Information

---

What is your year of birth? (e.g., 1990)

Choose one or more races that you consider yourself to be:

- |  |   |
|--|---|
| <input type="checkbox"/> White                     | <input type="checkbox"/> Hispanic or Latino |
| <input type="checkbox"/> Asian                     | <input type="checkbox"/> Other              |
| <input type="checkbox"/> Black or African American |   |
- 

What is your gender?

- ☐ Male
- ☐ Female
- ☐ Other
- 

What is the highest level of school you have completed or the highest degree you have received?

- ☐ Less than High School
- ☐ High School
- ☐ College
- ☐ Graduate or Professional (JD, MD)

---

## Ending

---

This is the end of the survey. Thank you for your valuable time.

To obtain your payment, please input your unique ID below to MTurk.

Here is your unique ID: \${e://Field/Random%20ID}. Copy this value to paste into MTurk.

When you have copied this ID, please click the Submit button to submit your answers.

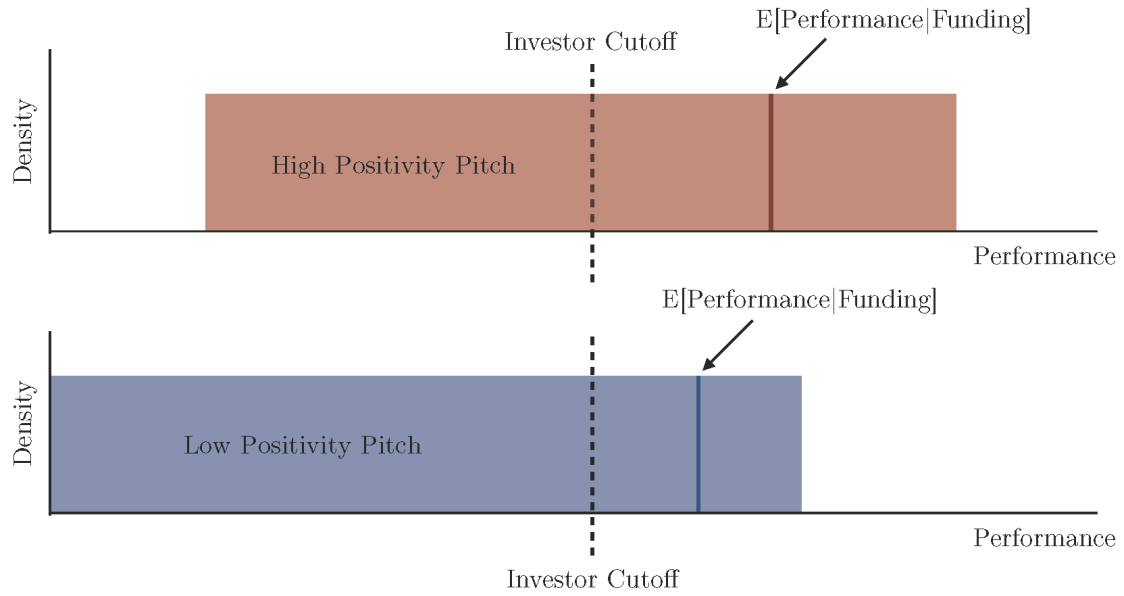
Powered by Qualtrics

## E. Appendix: Performance Analysis and Sources of Bias

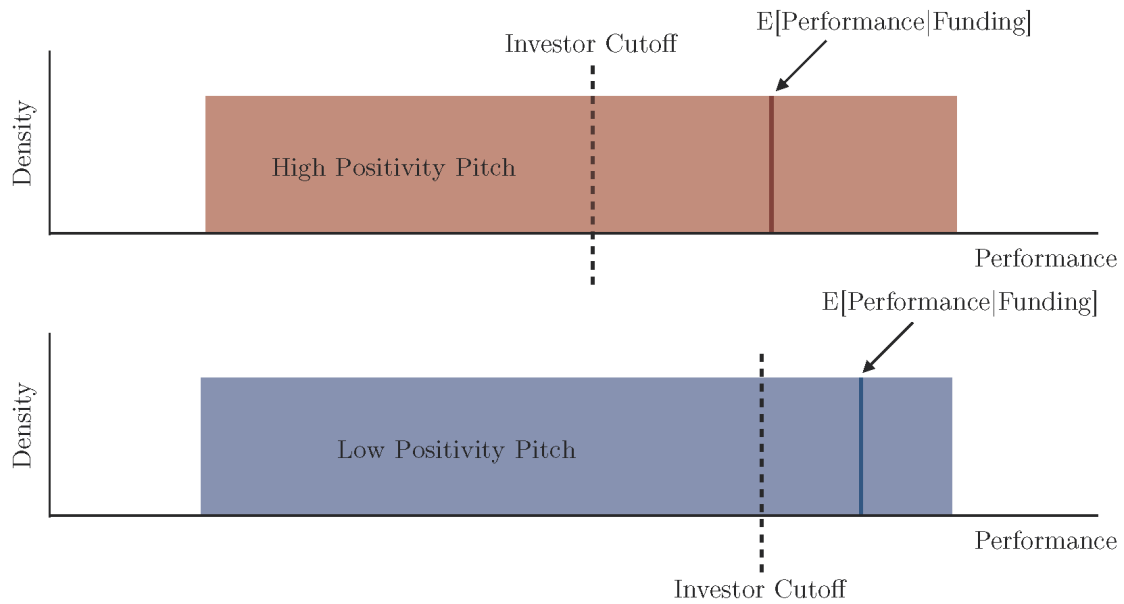
This appendix presents a simple conceptual framework, visualized in Figure [IA.7](#), to illustrate how pitch deliveries could introduce investment bias that then leads to poorer startup performance. Panel (a), presenting the no-bias scenario, shows hypothetical performance/quality distributions for startups that an investor may be considering funding. Separate overlapping distributions are assumed for startups with high- versus low-positivity pitches. The distributions shown are identical, except that the high-positivity distribution is shifted to the right of the low-positivity distribution. In other words, the high-positivity teams first-order stochastically dominates the low-positivity distribution. We assume the investor funds startups according to a simple cutoff rule, offering funding to all startups above a certain threshold. Since the investor is unbiased, he or she applies the same cutoff rule to all startups, regardless of the pitch positivity. In this case, because the high-positivity distribution first-order stochastically dominates the low-positivity distribution, the investor will invest in startups with high-positivity pitches with greater probability. In addition, expected performance, conditional on funding, will be higher for high-positivity startups.

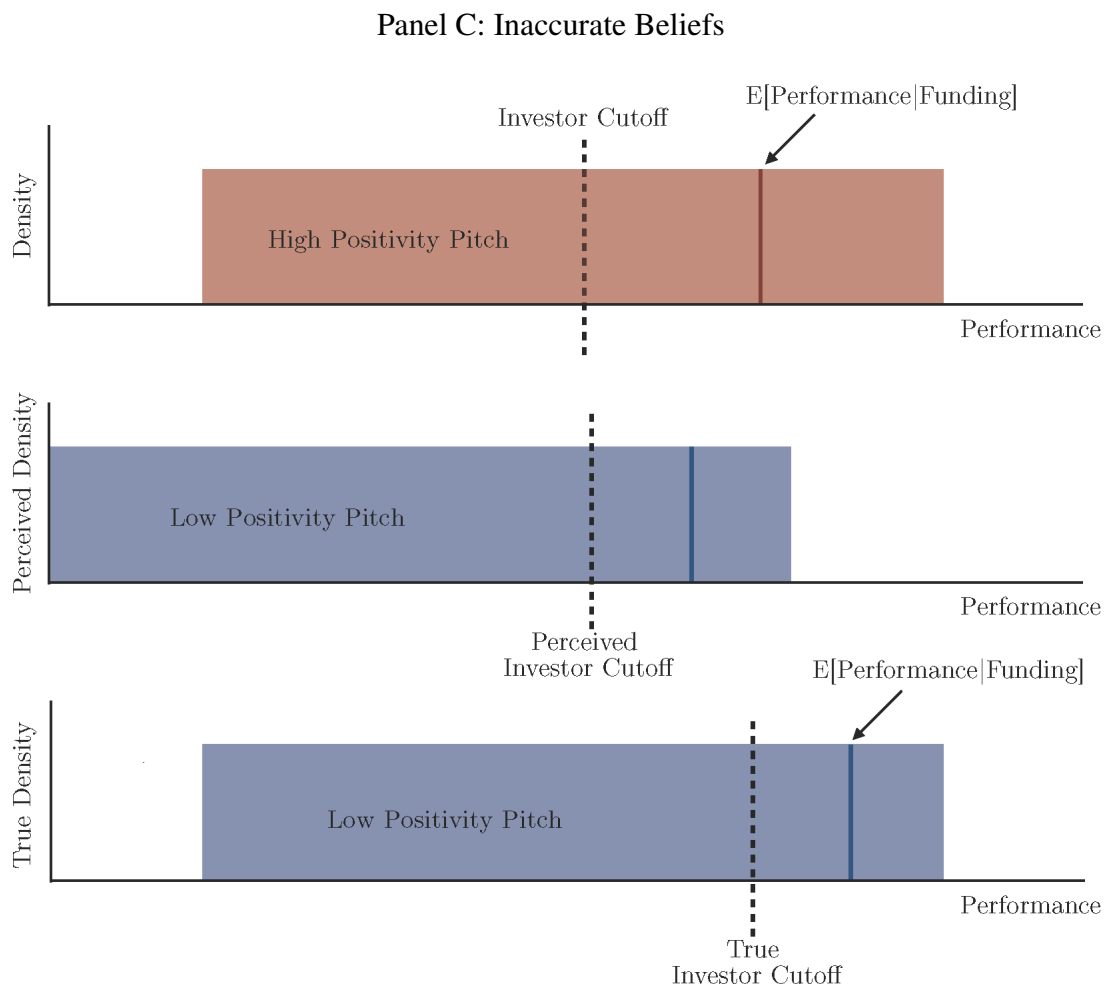
In contrast, if investors are biased, either due to a taste-based channel or inaccurate beliefs, it is possible that high-positivity startups may underperform. Figure [IA.7](#) Panel (b) illustrates taste-based bias. In the example, the performance distributions of high- and low-positivity teams are assumed to be the same. The investor continues to derive utility from startup performance. But she or he now also derives disutility from investing in startups with low positivity pitches—as a result, the investor sets a higher cutoff for them. With a taste-based channel, the investor will again fund founders with more positive pitches with greater probability. However, now expected performance, conditional on funding, will be lower for these investments. Figure [IA.7](#) Panel (c) illustrates the case of inaccurate beliefs. Inaccurate beliefs imply a gap between the investor’s perceived performance distribution for low-positivity (or high-positivity) startups and the true performance distribution. In the example shown, the investor acts exactly like an investor with no bias according to the investor’s *perceived* performance distribution. Inaccurate beliefs can also lead investors to fund founders of high-positivity with greater probability while having lower (true) expected performance for those investments.

Panel A: No Bias



Panel B: Taste-Based Bias





**Figure IA.7. Startup Performance Under Different Investment Models.** These figures present hypothetical startup performance distributions combined with investor decision rules. Panel A considers the situation where the investors have no bias and startups with low-positivity pitches underperform high-positivity startups. Investors use the same performance cutoff rule (the vertical dashed line) and the solid vertical lines represent the expected performance conditional on the funding decision. Panel B considers the situation where investors exhibit taste-based bias and founders of both high- and low-positivity have the same performance distribution. The taste-based bias leads investors to have a higher cutoff rule (the vertical dashed line) for low-positivity startups. This, in turn, leads to higher performance outcomes conditional on funding. Panel C presents the situation where investors have inaccurate beliefs about startups with different pitch features. The low-positivity startups' distribution is shifted to the left because of the miscalibration, which has the effect of increasing the expected performance conditional on funding.



## **F. Appendix Figures and Tables**

**Table IA.9**  
**Features in Pitch Delivery and Investment Decisions: MSFT Azure**

Logit regressions, marginal effect,  $N = 1,139$ . The analysis is obtained using the following model:

$$I(Invested) = \alpha + \beta \cdot X + \delta_{FE} + \varepsilon.$$

$I(Invested)$  takes a value of one if the startup team was chosen by the accelerator, and zero otherwise. All pitch feature variables are standardized into a zero-mean variable with a standard deviation of one. Visual variables are constructed by Microsoft Azure APIs. Vocal and verbal variables are identical to those in Table IV. Pitch Factor is constructed by the same method as in Table IV. All regressions include Accelerator FE. Control variables include founders' education background (whether they have a Master's or a PhD degree, whether they attended an elite university, defined as the U.S. News & World Report's Top 10), founders' prior work experience (whether they have prior entrepreneurship experience, whether they ever held a senior position in prior employment), team size, video resolution, and informational content. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

Dependent Var: $I(Invested)$	w/ Content Controls		w/ Content and Startup/Team Controls	
	Marginal Effect	S.E.	Pseudo $R^2$	Pseudo $R^2$
Pitch-Factor	0.028***	(0.007)	0.191	0.251
<i>Visual (Facial)</i>				
Visual-Positive	0.012**	(0.006)	0.176	0.239
Visual-Negative	-0.013*	(0.007)	0.176	0.239
Visual-Beauty	0.015**	(0.006)	0.178	0.242
<i>Vocal (Audio)</i>				
Vocal-Positive	0.009**	(0.005)	0.174	0.239
Vocal-Negative	-0.045***	(0.016)	0.183	0.248
Vocal-Arousal	0.023***	(0.009)	0.184	0.245
Vocal-Valence	0.023***	(0.006)	0.185	0.246
<i>Verbal (Text)</i>				
Verbal-Positive	-0.010	(0.009)	0.174	0.239
Verbal-Negative	-0.026***	(0.007)	0.186	0.246
Verbal-Warmth	0.026***	(0.008)	0.190	0.256
Verbal-Ability	-0.049***	(0.009)	0.243	0.298

**Table IA.10**  
**Features in Pitches and Investment Decisions: Robustness Checks**

Logit regressions, marginal effect. The analysis is obtained using the following model:

$$I(Invested) = \alpha + \beta \cdot X + \delta_{FE} + \varepsilon.$$

$I(Invested)$  takes a value of one if the startup team was chosen by the accelerator, and zero otherwise. All pitch feature variables are standardized into a zero-mean variable with a standard deviation of one. All variables are identical to those in Table IV. Control variables include founders' education background (whether they have a Master's or a PhD degree, whether they attended an elite university, defined as the U.S. News & World Report's Top 10), founders' prior work experience (whether they have prior entrepreneurship experience, whether they ever held a senior position in prior employment), team size, and video resolution. Standard errors clustered at the accelerator-year level are displayed in parentheses. \*\*\*, \*\*, and \* indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
	Dependent Var: $I(Invested)$			
Pitch Factor	0.028*** (0.010)	0.023** (0.010)	0.014** (0.007)	0.026*** (0.008)
Observations	1,139	1,139	1,139	1,139
Specification	OLS	OLS	Logit	Logit
$R^2$ /Pseudo $R^2$	0.151	0.201	0.440	0.312
Startup/Team Controls		Y	Y	Y
Content Controls		Y	Y	Y
Accelerator FE	Y	Y		Y
Accelerator-Year FE			Y	
Industry FE				Y

**Table IA.11**  
**Gender Breakdown by Industry**

This table provides industry (GICS) distributions of collected videos across different team gender compositions.

	Men-Only	Women-Only	Mixed-Gender
Communication Service	4.83	7.10	4.81
Consumer Discretionary	20.57	21.94	15.19
Consumer Staples	2.50	6.13	2.59
Energy	0.36	0.65	0.00
Financials	5.19	5.16	4.07
Health Care	6.62	8.06	10.00
Industrials	7.69	8.39	9.63
Information Technology	48.12	37.42	50.00
Materials	0.18	0.65	0.00
Real Estate	1.97	0.97	1.48
Unclear	1.97	3.55	2.22
Total Observation	559	310	270
Total %	100.00	100.00	100.00

## G. Experiment: Summary Statistics and Sample

**Table IA.11**  
**Summary Statistics of Subjects in Experiments**

This table provides descriptive statistics of demographic information of subjects in our experiment sample. The demographic information is collected during the experiment. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles.

	N	Mean	STD	25%	50%	75%
Age	102	28.35	3.31	25.00	28.00	31.00
Man	102	0.60	0.49	0.00	1.00	1.00
Woman	102	0.40	0.49	0.00	0.00	1.00
White	102	0.45	0.50	0.00	0.00	1.00
Black or African American	102	0.03	0.17	0.00	0.00	0.00
Asian	102	0.42	0.50	0.00	0.00	1.00
Hispanic or Latino	102	0.05	0.22	0.00	0.00	0.00
Mixed Race	102	0.03	0.17	0.00	0.00	0.00
Other Race	102	0.02	0.14	0.00	0.00	0.00

**Table IA.12**  
**Summary Statistics of Unstandardized Features in Experiments**

This table provides descriptive statistics of pitch features. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles. Variables are categorized into vocal, video, and verbal.

	N	Mean	STD	25%	50%	75%
<i>Visual (Facial)</i>						
Visual-Positive	62	0.18	0.17	0.06	0.13	0.30
Visual-Negative	62	0.17	0.18	0.05	0.10	0.24
Visual-Beauty	62	0.59	0.09	0.52	0.60	0.64
<i>Vocal (Audio)</i>						
Vocal-Positive	62	0.08	0.04	0.04	0.07	0.09
Vocal-Negative	62	0.02	0.01	0.01	0.01	0.02
Vocal-Arousal	62	0.35	0.35	0.09	0.23	0.67
Vocal-Valence	62	0.28	0.26	0.08	0.22	0.49
<i>Verbal (Text)</i>						
Verbal-Positive	62	0.02	0.01	0.01	0.01	0.02
Verbal-Negative	62	0.01	0.01	0.00	0.01	0.02
Verbal-Warmth	62	0.02	0.02	0.01	0.02	0.02
Verbal-Ability	62	0.03	0.03	0.01	0.03	0.04

**Table IA.13**  
**Summary Statistics of Video Pitches in Experiments**

This table provides descriptive statistics of basic information of the pitch videos. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles.

	N	Mean	STD	25%	50%	75%
Duration (second)	62	61.76	4.88	58.00	61.00	66.00
Video Size (MB)	62	12.79	10.22	4.55	9.10	17.06
Number of Words	62	174.74	39.34	149.00	176.00	199.00
Number of Sentences	62	11.65	3.53	9.00	11.50	13.00
Number of Views	62	2,742.06	14,558.83	65.00	149.50	327.00
Number of Likes	62	3.03	7.53	0.00	0.00	2.00
Number of Dislikes	62	0.24	0.97	0.00	0.00	0.00

**Table IA.14**  
**Summary Statistics of Startups in Experiments**  
**(as of July 2019)**

This table provides descriptive statistics of characteristics of startups all measured as of July 2019 from Crunchbase and PitchBook. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles.

	N	Mean	STD	25%	50%	75%
Invested by Accelerator	62	0.44	0.50	0.00	0.00	1.00
Firm Age	62	3.44	1.71	2.00	3.00	5.00
Number of Employees	32	26.56	70.81	5.00	5.00	30.00
Startup Alive	32	0.91	0.30	1.00	1.00	1.00
Raised VC	32	0.53	0.51	0.00	1.00	1.00
Total Funding Amount (\$000)	32	12,685	47,022	0.00	148	2,700

**Table IA.15**  
**Summary Statistics of Teams in Experiments**

This table provides descriptive statistics of the startup teams. Team member background information is collected from LinkedIn. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles.

	N	Mean	STD	25%	50%	75%
Number of People	62	2.10	1.20	1.00	2.00	3.00
Single-Member	62	0.34	0.48	0.00	0.00	1.00
Multi-Member	62	0.66	0.48	0.00	1.00	1.00
Men-Only	62	0.52	0.50	0.00	1.00	1.00
Women-Only	62	0.32	0.47	0.00	0.00	1.00
Mixed Gender	62	0.16	0.37	0.00	0.00	0.00
Has LinkedIn	62	0.73	0.45	0.00	1.00	1.00
Prior Senior Position	45	0.82	0.39	1.00	1.00	1.00
Prior Startup Experience	45	0.58	0.50	0.00	1.00	1.00
Elite University	45	0.13	0.34	0.00	0.00	0.00
Master Degree	45	0.33	0.48	0.00	0.00	1.00
PhD Degree	45	0.13	0.34	0.00	0.00	0.00

**Table IA.16**  
**Summary Statistics of Beliefs and Investment Decisions in Experiments**

This table provides descriptive statistics of beliefs and investment decisions elicited in the experiment. For each variable, we report the number of observations, mean, standard deviation, and 25th, 50th, and 75th percentiles.

	N	Mean	STD	25%	50%	75%
<i>Belief (<math>\mu</math>)</i>						
Baseline $P(\textit{invested})$	952	0.20	0.17	0.08	0.15	0.29
Baseline $P(\textit{alive} \textit{invested})$	952	0.25	0.15	0.12	0.20	0.32
$I(\textit{invested})$	952	0.46	0.50	0.00	0.00	1.00
$P(\textit{alive} \textit{invested})$	952	0.31	0.23	0.14	0.26	0.45
$P(\textit{alive} \textit{not invested})$	952	0.17	0.18	0.05	0.10	0.24
$P(\textit{success} \textit{invested})$	952	0.13	0.18	0.02	0.05	0.17
<i>Precision of Belief (<math>\sigma</math>)</i>						
Baseline $P(\textit{invested})$	952	3.30	0.79	3.00	3.00	4.00
Baseline $P(\textit{alive} \textit{invested})$	952	3.24	0.69	3.00	3.00	4.00
$I(\textit{invested})$	952	2.60	0.90	2.00	3.00	3.00
$P(\textit{alive} \textit{invested})$	952	2.74	0.85	2.00	3.00	3.00
$P(\textit{alive} \textit{not invested})$	952	2.74	0.86	2.00	3.00	3.00
$P(\textit{success} \textit{invested})$	952	2.73	0.88	2.00	3.00	3.00

## Consent Form

---

Hi, this is a survey designed by the research team of Song Ma (Assistant Professor of Finance at Yale School of Management). We are conducting research to examine the relation between entrepreneurs' performance in video pitching and their outcomes in obtaining venture investment.

We are inviting you to participate in this study by completing this short survey. This survey will take you around 20 minutes. The results of the survey will be used for research purposes only. All of your responses will be held in confidence.

**This survey is also a required assignment of MGT 897 - Entrepreneurial Finance.** You will get a base point of 5 as long as you finished this survey. In addition to your base point, we will award you bonus credits. The bonus credit (up to 3 points) is determined by how well you did in the survey (e.g., you choose to invest in an entrepreneur team that later became more successful.)

Would you like to participate in the study?

- ☐ Yes
- ☐ No

---

## Basic Information Section

---

What is your Yale NetID?



---

---

What is your year of birth? (e.g., 1990)

---

Which program do you currently enroll at Yale University?

- ☐ Undergraduate
  - ☐ Master at Yale SOM (e.g., MBA, EMBA, MAM, and MMS)
  - ☐ PhD
  - ☐ Other
- 

Which year are you in the current program at Yale University?

- ☐ First year
  - ☐ Second year
  - ☐ Third year and above
- 

Choose one or more races that you consider yourself to be:

- |  |   |
|--|---|
| <input type="checkbox"/> White                     | <input type="checkbox"/> Hispanic or Latino |
| <input type="checkbox"/> Asian                     | <input type="checkbox"/> Other              |
| <input type="checkbox"/> Black or African American |   |
- 

What is your gender?

- ☐ Male
- ☐ Female

☐ Other

---


Which of the following categories best describes your previous occupation? (Choose at least one and no more than four)

- |  |   |
|--|---|
| <input type="checkbox"/> Student                         | <input type="checkbox"/> Entrepreneur                       |
| <input type="checkbox"/> Asset Management and Banking    | <input type="checkbox"/> Technology                         |
| <input type="checkbox"/> Consulting                      | <input type="checkbox"/> Venture Capital and Private Equity |
| <input type="checkbox"/> Education                       | <input type="checkbox"/> No Full-time Work Experience       |
| <input type="checkbox"/> Energy/Healthcare/Manufacturing | <input type="checkbox"/> Other                              |
- 

### Benchmark Belief Section

---

**On average, what percentage** of startups do you think can successfully raise Series A financing from VC conditional on trying?



How confident are you with your answers to the question about the probability of obtaining the investment that you were just asked?

- ☐ Extremely confident
- ☐ Very confident
- ☐ Somewhat confident
- ☐ Not very confident
- ☐ Not at all confident
-

If a startup has already been invested by a venture capital, what do you think is the **average** successful rate of a startup to survive in the following three years?



How confident are you with your answers to the question about the surviving probability that you were just asked?

- ☐ Extremely confident
- ☐ Very confident
- ☐ Somewhat confident
- ☐ Not very confident
- ☐ Not at all confident

### Video Pitch Experiment Introduction

Now, imagine that you are a venture investor. **You are going to decide whether to invest in a given startup after watching its one-minute video pitch.** If you decide to invest in this startup, the contract will be the same – you will invest \$150K in this startup team for 7% share of the company.

In the following part of the survey, you are going to watch **10** video pitches and decide whether to invest in these startups.

***Note: The submission button for each page will appear only after the video is watched and all questions are answered. If the submission button does not***

***appear even after all questions are answered, please wait several seconds and do not reload the web page. (Reloading will only reset the your answers.)***

---

### **Video Pitch IY3hoi1eizM (Example)**

---

**These page timer metrics will not be displayed to the recipient.**

First Click: *0 seconds*

Last Click: *0 seconds*

Page Submit: *0 seconds*

Click Count: *0 clicks*

---

Y-Combinator Application Video - 1min



**Please watch the video. All survey questions are related to this video.  
(The submission button will appear after the video is played and questions are answered.)**

---


If you were an investor, are you willing to invest \$150K in this startup team for 7% share of the company?

- ☐ Yes
  - ☐ No
- 

How confident are you with your answers to the question about the investment decision that you were just asked?

- ☐ Extremely confident
  - ☐ Very confident
  - ☐ Somewhat confident
  - ☐ Not very confident
  - ☐ Not at all confident
- 


If this startup **was able to raise Series A financing from VC**, what is the probability that you think this startup will still be alive (including being acquired) three years later?



How confident are you with your answers to the question about the surviving probability that you were just asked?

- ☐ Extremely confident
  - ☐ Very confident
  - ☐ Somewhat confident
  - ☐ Not very confident
  - ☐ Not at all confident
-


If this startup **was not chosen by a VC**, what is the probability that you think this startup will still be alive (including being acquired) three years later?



How confident are you with your answers to the question about the surviving probability that your were just asked?

- ☐ Extremely confident
- ☐ Very confident
- ☐ Somewhat confident
- ☐ Not very confident
- ☐ Not at all confident

What is the probability that you think this startup will become a huge success (e.g., a Unicorn)?



How confident are you with your answers to the question about the huge success probability that your were just asked?

- ☐ Extremely confident
- ☐ Very confident
- ☐ Somewhat confident

- ☐ Not very confident
- ☐ Not at all confident
- 

What are the most important factors in your decision of whether to invest in **this startup?**

	Extremely important	Very important	Somewhat important	Not very important	Not at all important
Team's pitching traits (e.g., facial expression, passionate voice, beauty)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Team's general ability	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Team's general sociability	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Products, business model, industry, and market	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Team's previous industry experience	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Team's previous entrepreneurial experience	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Team's education background	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

---

## Ending

---

This is the end of the survey. Thanks for your valuable time.

If you have any additional comments about this survey, please provide them below.  
(Optional)

Powered by Qualtrics