The Informational Role of Imagery in Financial Decision Making: A New Approach*

Susan Gans[d], Tavy Ronen[a], Tawei (David) Wang[b], Mi (Jamie) Zhou[c]

[a] *Rutgers Business School*
[b] *DePaul University*
[c] *Virginia Commonwealth University*
[d] *West Chester University*

Abstract

What is the incremental information in imagery? This paper introduces the concept of visual attention and proposes a new methodology to identify the impact of images and of their information content on financial decision making. We introduce two sets of criteria for analyzing the effects of images - visual clarity and information reinforcement, which together inform individual investors. Our flexible methodology can be seamlessly integrated with modern information processing techniques, from simple keyword matching to advanced semantic analysis or topic modeling. Using images as data and applying the methodology to a hand-collected data set of equity crowdfunding campaigns, we show that the existence of perceptible images is significantly associated with greater investment, and that visual clarity measures are found to incrementally affect investor decision making after controlling for textual sentiment and financial information. Specifically, more vibrant images with more focal points are associated with higher funding. Additionally, information reinforcement is found to be a key determinant of (limited) visual attention. Data collected in a human participant sample supported these findings. Finally, the methodology we develop in this paper provides a new framework within which to discuss (limited) attention, information, news, and investor reactions in today's digital and visual age.

Keywords: image quality, visual attention, limited attention, informational efficiency, textual analysis, machine learning, novel methodology equity crowdfunding, classification
JEL Classification:G0

*This paper has benefited from comments by Azi Ben-Rephael, Hsiu-Lang Chen, Serdar Dinc, Mark Fedenia, Susan Gans, Joel Hasbrouck, Sofia Johan, Marios Panayides, Asheq Rahman, Joshua Ronen, Tom Scott, Chris van Staden, Charles Trzcinka, Paul Wells, Ziwei Zhao, Weinan Wheng, the 2018 Decision Science Meetings, the Rutgers Business School Finance Department, the RBS Center for Business of Fashion Symposium, the 2020 FMA Meetings, PBFEAM 2021 Meetings, and the 2021 EFA Meetings. We would like to thank West Chester University of Pennsylvania for use of their laboratory facilities, and Mr. Derrick Stahl for executing the human participant experiment. Tawei Wang received financial support from DePaul University.

This draft: May 2021 First draft: August 2019

1. Introduction

> "In photography there is a reality so subtle that it becomes more real than reality"
>
> –*Alfred Stieglitz*

Photographers and artists have long been privy to the secret powers of imagery on the human mind. Today, our lives and economy have become overtaken by an endless sea of snapshots, screenshots, filters, memes, emojis, and selfies. Everywhere we turn, we are bombarded with visual pleas, fighting for our attention. In psychology, terms like "mere exposure", "arousal dynamics" and "perceptual and conceptual fluency" permeate the field.[1] Top photographers engage consumers using Gestalt principles, and social media influencers battle it out on Instagram.[2]

Given the pervasive use of imagery in today's rapidly evolving and visual world, the dearth of attention it receives in the financial economics literature is surprising.[3] In this paper, we argue that imagery is a missing factor in the current literature on attention and information analysis and highlight the importance of pictorial representation on financial decision making. Unlike behavioral finance studies, which examine the irrational or subjective reaction of investors to psychological or other triggers, our approach focuses on the objective, or rational reactions of individual investors to concrete and measurable elements of pictorial representations. We introduce the term visual attention and propose an objective set of metrics based on two sets of criteria for analyzing the effects of images – visual clarity and information content/reinforcement, which together inform individual investors.

The first dimension we propose, information content/reinforcement, involves isolating the

[1] See Palmer et al. (2013) for an analysis on aesthetic preferences in experimental psychology.

[2] See Wagemans et al. (2012) for a review on Gestalt Psychology in Visual Perception and Schroeder (2008) for a discussion on issues of visual images as they pertain to brands.

[3] While studies in advertising, consumer behavior and marketing have looked at the value of advertising with images, the issues and methodologies in this study have not been analyzed. Further, those papers do not examine the effect on financial investment in companies, rather on consumer behavior. Other papers focused on behavior of investors do not examine the impact of specific images. For example, while Liaukonyt et al. (2018) find that within 15 minutes of national TV advertisements, EDGAR queries are increased, the content of the visual images is not addressed.

information content of pictorial representations from that in contemporaneously displayed text, thereby overcoming a common criticism of extant studies analyzing short term news or media effects on financial outcomes in their inability to control for potentially confounding information that is concurrently released. Our approach involves multiple stages. First, we propose using a novel machine learning based classification scheme which in part relies on Google Vision (henceforth GV) label identification, to ascertain whether images contain informative content (provide material information which can be synthesized by investors) by examining GV image labels.[4] Next, we compare these labels to words and phrases in the accompanying text, and determine whether the abstracted information is 'reinforcing' or 'additive' (to accompanying textual information). This last classification is key in allowing us to distinguish between reactions to the different types of information and to disentangle the impact of imagery. Our methodology is flexible in the degree to which textual keywords are identified, the number of reinforcement categories considered, the thresholds required for categorization, and can be seamlessly integrated with various modern information processing techniques, from simple keyword matching to advanced semantic analysis or topic modeling.

The second dimension we propose, 'visual clarity' is comprised of three metrics which have been used in the visual saliency and psychology literatures, in different contexts: color intensity, singularity, and resolution. We posit that resolution and color intensity (how colorful images are) should positively affect investors, and that the effect of singularity, the degree to which an image focuses on one or few objects, would likely be negative. Finally, we consider potential interactions between these two dichotomies (reinforcement and clarity).

A natural testing ground for our somewhat novel theories exists in the world of equity crowdfunding. Since crowdfunding is rich in both textual and visual representation, and since it typically attracts less sophisticated (but web-savvy) investors, we expect that individual investors would rely in part on the visual elements of the campaigns to glean information regarding the quality of the projects or companies, beyond just the text-based descriptions and

[4](Google Vision API, Google (2019))

other firm indicators. Further, we argue that since crowdfunding tends to attract retail investors with limited contribution amounts, the amount of effort placed on campaign valuation may be limited as well, rendering easier-to-process information, such as visual cues, non-trivial. Finally, unlike projects of public companies, the information environment of these projects is relatively controlled, with limited (at best) access to information beyond that listed on the platform. Our hand-collected dataset comprises two popular equity crowdfunding platforms: EquityNet and Crowdcube, which represent different spectrums of visual representation. The EquityNet platform, which does not allow companies to post images other than logos, is a controlled visual environment which is most conducive to isolating the effect of including an image along with textual information, constituting a precursor to examining our visual attention methodology in the more visually sophisticated Crowdcube platform. We find strong evidence of an 'image effect' based on logos in the EquityNet sample and show that for a large cross section of campaigns, the existence of a perceptible image is significantly associated with greater investment amounts (by up to 65%) after controlling for financial information, textual sentiment (in accompanying text), and other firm characteristics.

The more visually complex Crowdcube equity crowdfunding platform (which in addition to logos also displays rich background images) provides a natural laboratory environment within which to test our visual attention methodology and to examine the impact of our two sets of criteria, clarity and information content/reinforcement, on financial investment levels. We begin by reaffirming the existence of an image effect: companies that include background images in addition to logo images raise on average 13% (or roughly £83,322) more than those that do not. This result suggests that the strong evidence of an image effect found in the Equitynet sample is not an artifact of logo-specific features. Next, we find our visual clarity measures to be significantly associated with investment amounts. Specifically, after controlling for sentiment and financial variables, we find that background image intensity (singularity) is positively (negatively) and significantly associated with the amount of funds raised. The economic significance of visual clarity can be substantial. For example, based on the average dollar amount of funds raised per Crowdcube campaign (roughly £640,935) during our sample period), a unit decrease in image singularity (interpreted as an increase in the number of focal points) is associated with on average, a £37,302 larger project investment amount, and a unit increase in image intensity (vibrancy or dolorfulness) is associated with on average, a £103,831 larger investment.

Importantly, however, our methodology, which allows us to extract the incremental information content of images (above and beyond that of accompanying text), allows us to discern that these results are driven by those campaigns for which the background images are both informative and informationally additive (not reinforcing). That is, not only are images with 'additive' information content (to that in accompanying textual information) determined to be more valuable to investors than those which are merely 'reinforcing,' but clarity metrics are only significant for the additive group. We interpret these results as consistent with our simple visual attention story: when images contain information content that is not merely redundant to that in the accompanying text, they may receive more attention, perhaps in an attempt to resolve information uncertainty. And, those images that are more vibrant and have more (potentially informative) focal points may be those that help convey new information to investors.

We conclude that the informational reinforcement/additivity dimension of visual attention is a critical factor in determining the value of imagery and that while clarity can contribute to the financial success of a campaign, it is the additivity of the objective information content portrayed in the visual representations that is of greatest import. These results may provide practical implications for entrepreneurs as well as for larger corporations seeking capital or investor engagement. Companies may find it profitable to expend greater resources on imagery to ensure that it is both high in clarity and in information content, thereby capitalizing on investor sensitivity to visual attention.

The research design in this study is complemented by the identification of three choice testing grounds: EquityNet, Crowdcube, and field experiments. The EquityNet platform is most suitable to examine the image (logo) effect due to its controlled environment, in which only textual information and a logo are visible to investors. We then move to our main testing ground, Crowdcube, in which investors see mainly a large background (cover) image, a logo, and textual information. Other images and videos may exist, but unlike on popular product crowdfunding platforms, they are visually secondary, often visible only at the fixed frame of the page and are not embedded in project description.[5] Since our focus is on the value of images, and not of the content (or mood) in videos (see Hu and Ma (2020) for an analysis of how videos affect investors in product crowdfund pitches), these simpler platforms comprise more controlled environments for us to examine visual attention arising from still images. Finally, we supplement our study with data collected from human participants in a laboratory experiment. The data from our human participants endorses our analytic approach to images, confirming that the same image qualities that our analyses identify are also perceived by people as they make investment choices.

Our study straddles and contributes to several streams of literature. First, in contrast to a stream of literature examining the peer-to-peer lending market which links attractiveness to investor behavior such as Duarte et al. (2012), Pope and Sydnor (2011), and Ravina (2019), we do not attempt to capture the emotional effects of image appeal, rather we measure the impact of identifiable elements of concrete information.[6] Our focus is on disentangling the information conveyed by images and showing how it may add to textual information that may in turn affect investors' investment decisions.

The approach taken in this paper also bears on the areas of attention (and limited attention) and news. Recent work in this area has examined the impact of the timing and method of dissemination on investor attention, such as Fedyk (2018), who finds that the placement

[5]See for example Kickstarter and Indiegogo, in which the main visual prompt for investors is a cover video. On these sites, images and videos may also be embedded elsewhere within textual descriptions.

[6]While images may certainly also convey subjective or behavioral effects not easily captured by economists, the methodology proposed in this paper focuses on an objective measurement. Behavioral studies documenting images effects on consumer decision making include Jarvenpaa (1990), Mandel and Johnson (2002), and Ackert et al.. Those studies examine emotional responses of consumers in individual purchasing decision and do not consider financial investor behavior.

of firms' financial information on the Bloomberg Terminal is more important than the actual content, Da et al. (2014), who show that continuous bits of information attract less attention than larger, discrete chunks of news, and Dellavigna and Pollet (2009) who find that stock price responses are weaker when earnings announcements are on Friday. Other studies have linked advertising to financial market outcomes.[7] Notably, this body of literature has not to date considered the impact of imagery on investor attention or on asset prices.[8] We believe that (limited) visual attention, a concept we introduce here, is a natural extension of this literature, particularly in light of studies such as Hirshleifer and Teoh (2003) who argue that investors are more likely to consider easy-processed information than more complicated items and that the form of presentation matters.

Finally, our work directly contributes to the literature on hard and soft information. These distinct types of information have been studied in both accounting and finance. In the accounting literature, hard disclosures are considered more credible and informative than soft disclosures because they are more verifiable, precise and objective (for example, Plumlee et al. (2015)). In the finance literature, examples of hard information can include changes in the market index or stock price, whereas the quality of a new management team or a borrower's ability to repay loans may be considered to be soft. In the crowdfunding context, Knyazeva and Ivanov (2017) show that soft information plays a significant role in the US securities-based crowdfunding market (compared to hard information). They find that an issuer's ability to amass a significant social media following is viewed as a positive signal of issuer quality, po-

[7]See also Chemmanur and Yan (2009) and Lou (2014) who demonstrate that advertising increases investor attention and is associated with larger stock returns, Liaukonyt et al. (2018) find that a firm's TV ad spurs EDGAR searches on the firm, Mayer (2019) who find that ads during college football games create price pressure in the stock market, Madsen and Niessner (2019) who find that print ads (mainly in business publications) cause transitory investor attention spikes, Solomon (2012) who suggests that firms hiring investor relations firms experience stronger market reactions around new announcements, and Focke et al. who show that advertising can create investor attention but that the economic significance is small and raise concerns regarding causality, and Ben-Rephael et al. (2017), who differentiate between retail and institutional attention.

[8]Two interesting industry articles focusing on small businesses and fashion companies, respectively (Antonelli (2016), Walker (2018)), discuss the importance of first impressions, typography and graphic elements in logos. However, the advertising literature has been subject to criticism in terms of the researcher's ability to determine whether better firms can advertise more, or whether advertising affects financial metrics.

tentially capturing a favorable outlook for market interest in the issuer's product or service and the recognition of the issuer's brand.

The methodology we develop in this paper can hopefully provide the seeds for a new framework within which to discuss news, information and investor reactions in today's digital and visual age. Recent studies focusing on decoding textual analysis in news (financial disclosures, internet searches, news stories and social media) and condensing it into numerical indexes have opened new frontiers in terms of understanding how reported information is incorporated by investors, but are limited how well they can determine how the human mind assimilates the information in text.[9] While this paper faces similar challenges, the goal is not the same. The textual analysis literature seems to strive towards perfecting the ability to harden soft information.[10] In this study, we consider pictorial information to be neither hard nor soft, and use objective metrics to decode, categorize it, and demonstrate how the incremental information conveyed by images can contribute to the information set investors use in making financial decisions.

The remainder of the paper is organized as follows. In Section 2, we provide a brief background description on the context for our empirical application, equity crowdfunding and we also provide an overview on the visual saliency literature. Section 3 presents our methodology and provides illustrative examples of our textual and imagery decoding approach. Section 4 describes the data used in our empirical application, Section 5 presents the results and Section 6 concludes.

[9]See Liberti and Petersen (2019)for a survey and summary of this. Li (2008), Tetlock (2010), Loughran and Mcdonald (2011) Dougal et al. (2012), Huang et al. (2014), Loughran and McDonald (2014), Hoberg and Phillips (2009),Gentzkow et al. (2017), Giannini et al. (2011), among others for recent work in this area.

[10]Liberti and Petersen (2019).

2. Background on Visual Characteristics and Crowdfunding

2.1. Visual Characteristics and Imagery

Image characteristics have been discussed mainly in the image processing literature, and particularly in studies examining the aesthetic classification of images (see for example Desnoyer and Wettergreen (2010)). Studies have considered several features, such as "the Rule of Thirds" in photography, to determine the quality of an image (Datta et al. (2006)). Specific features, such as color, brightness, and focus are extracted and considered as classifiers. However, it is often challenging to determine the appropriate features that can be used as classifiers in different contexts, which has led to more generic local descriptors to characterize images through more sophisticated machine learning image processing algorithms (Marchesotti et al. (2015)).

While our study focuses on objective or rational measures, a preliminary, and therefore cursory, exploration of the psychology literature reveals that much work has been done to understand the qualities of images that are associated with aesthetic preferences and/or positive affect (these two measures are occasionally orthogonal). For example, positive affect can result from the completion of an incomplete image (Harris et al. (1972)) or from the direction in which image subjects gaze (Chen et al. (2018)). Rounded objects are generally preferred to angular ones (Silvia and Barona (2009)), and images of natural objects generate greater (aesthetic) agreement than those depiting man-made items (Vessel et al. (2018)). Researchers have also begun to identify the underlying neural pathways for these aesthetic/affective preferences (Belfi et al. (2019)).

In addition to the above psychology, algorithm and methodology related studies, marketing researchers have focused on the importance of visual messages on customer behavior. For example, more than three decades ago, Mitchell and Olson (1981) manipulated verbal versus visual cues in an experiment and showed that both produce attribute beliefs and termed attitude toward advertisement mediate brand attitudes. In addition, Luffarelli et al. (2019) study logo design and Zhao et al.(2009) examine the visualization of product information.

Larsen et al. (2004) propose a taxonomy that identifies the angle of vision, cutting rate and camera motion as ad system attributes. Miller and Kahn (2005) find that customers react favorably to unusual colors and product names, and Zhang et al. (2018) demonstrate that image quality can affect Airbnb booking volume. Hsiao et al. (2019) show that when designer brands' lookbooks include images from national brands, sales are boosted for small private label products, Malik et al. (2017) find that MBA student profiles photos on a professional site command a (subjective) beauty premium, and recent studies examining the peer-to-peer lending market link (subjective) attractiveness to investor behavior (Duarte et al. (2012), Pope and Sydnor (2011), and Ravina (2019)). Finally, researchers use images in the news to predict returns (Obaid and Pukthuanthong, 2020), employ visual characteristics of artwork to predict auction prices (Aubry et al., 2019), and study how visual salience affects investment decision (Bose et al., 2020). Additionally, Huang et al. (2020) and Hu and Ma (2020) investigate the relation between entrepreneurs' visual traits and venture capitalists' decision making.

2.2. Equity Crowdfunding

Equity crowdfunding has attracted much attention from entrepreneurs, the public and regulators in the past several years (Abate (2018), Cowley (2015)). Subsequent to the 2015 enactment of Securities and Exchange Commission's (SEC's) Regulation Crowdfunding ("Regulation CF"), equity crowdfunding has surged.[11] Abate (2018), for example, finds that within one year, more than 300 companies raised more than $30 million through equity crowdfunding campaigns. Unlike in non-equity crowdfunding campaigns, equity crowdfunding investors contribute money through the virtual platform in exchange for tangible interest, such as common stock. While equity crowdfunding campaigns are under the jurisdiction of securities laws

[11] `https://www.sec.gov/info/smallbus/secg/rccomplianceguide-051316.htm`. Equity crowdfunding, as distinct from other types of crowdfunding, rewards investors with equity in exchange for the funds. Since investment in these campaigns is therefore tantamount to the purchase of securities, several security laws have emerged to regulate such activity. Jumpstart Our Business Startups Act, Pub. L. No. 112-106, §§ 301–305, 126 Stat. 307, 315–323 (2012) (codified in 15 U.S.C. §§ 77a–77r, 78a–78o (2012)). The (JOBS) Act, which was enacted in 2016, allows for the general solicitation of accredited investors. In addition, the Capital Raising Online While Deterring Fraud and Unethical Non-Disclosures (CROWD-FUND) Act (part of the JOBS Act) allows for the participation of unaccredited investors (Guzik (2014)).

and regulations and need to meet certain compliance requirements, they offer small businesses and startups the opportunity to raise capital without having to meet stringent requirements. Typically low entry levels combined with caps on investment amounts help entrepreneurs, particularly smaller businesses, attract larger numbers of potential investors (Prive (2012), Wilmot (2018)), and the platform has been linked to social network contexts.

Despite the growing importance and popularity of equity crowdfunding, empirical studies are still limited. Studies that build on signaling theory, such as Ahlers et al. (2015) and Knyazeva and Ivanov (2017), have investigated factors that may be interpreted as effective signals that may lead to campaign success. Other studies have examined the success factors of these campaigns in a venture capital framework. For example, Mamonov and Malaga (2018) show that in selecting campaigns, investors focus on market and agency risks and that video narratives can affect equity crowdfunding success. Estrin and Khavul (2016) consider the quality of entrepreneurs and investors as factors that may affect the degree of contribution to the campaign. Scholars et al. (2016), on the other hand, demonstrate that funding history, media buzz, and Twitter presence can affect the success of campaigns. Yadav (2018) shows that from an entrepreneur's perspective, crowdfunding is instrumental in engaging new potential investors, which is at least as important as the venture finance component. Indeed, Vismara (2016) demonstrates that campaigns launched by entrepreneurs with higher levels of social capital are more likely to be successful, and Burtch, Ghose and Wattal (2013) discuss the social aspect of crowdfunding through which well liked campaigns are circulated and promoted amongst investors. Mollick (2014) shows that crowdfunding success appears to be linked to project quality (among other things), with projects signaling higher quality more likely to be funded. However, none of these papers link signaling to the quality or content of images.

3. Methodological Framework: A New Approach

The methodology we develop in this paper provides a new framework within which to discuss investor reactions to news and information by introducing a parallel dimension to textual and financial data: visual attention. We posit that two dimensions of objective criteria affect visual attention: clarity and information content/reinforcement, which together inform investors. We propose the use of objective metrics to calculate each of these dimensions, using visual saliency measures and machine learning classification schemes, respectively. Figure 1 depicts this framework.

———————Insert Figure 1 here———————

3.1. Visual Clarity

We define 'visual clarity' as the facility with which images can be observed, and identify clarity with three non-overlapping objective metrics, color intensity, singularity, and resolution, which are used in the visual saliency and psychology literatures, in different contexts.

3.1.1. Color Intensity

The first measure, color intensity, I, (also referred to as 'colorfulness' in the saliency literature), can be thought of as a main source of perceptual overall colorfulness of an image. As in Hasler et al. (2009), we first compute the following two metrics of opponent color space representation:

$$rg = R - G \tag{1}$$

$$yb = \frac{1}{2}(R + G) - B \tag{2}$$

where R is Red, G is Green, and B is Blue. We then calculate the standard deviation σ_{rgyb}, the mean μ_{rgyb} as:

$$\sigma_{rgyb} = \sqrt{\sigma_{rg}^2 + \sigma_{yb}^2} \tag{3}$$

$$\mu_{rgyb} = \sqrt{\mu_{rg}^2 + \mu_{yb}^2} \tag{4}$$

Finally we calculate the color intensity measure as:

$$I = \sigma_{rgyb} + 0.3 * \mu_{rgyb} \tag{5}$$

Figure 2 illustrates the intuition behind the intensity measure I. The image of the house in Panel A has a higher value of I (is more colorful) than the photograph of the brewing factory shown in Panel B.

——————Insert Figure 2 here——————

3.1.2. Image Singularity

Our second visual clarity metric, image singularity (S), captures the degree to which an image is focused on a single object (versus many). We construct S as the inverse of the image focus score, (σ_L^2), which in turn is calculated as the variance of L, where

$$L = \begin{pmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{pmatrix} \otimes M \tag{6}$$

and where M is the grey scale image. This focus score, (σ_L^2), is used in Pech-Pacheco et al. (2002) and essentially detects rapid intensity changes or edges in each image, such that when the matrix variance is large, it corresponds to multiple focal points in the image. Our measure, (S), is computed as the natural logarithm of the focus score (σ_L^2) times negative one (-1), and therefore increases as the number of focal points in an image decreases. The intuition behind S is simple: the larger the value of S, the less the number of focal points (objects) in the image. An illustrative example is provided in Figure 3. The image shown in Panel A has a higher singularity score than that in Panel B, since it is focused on one object (the toddler) unlike the four clear focal points (the beverage cans) depicted in Panel B.

——————Insert Figure 3 here——————

3.1.3. Image Resolution

Resolution, (R), is calculated as pixel width x height. This measure can be interpreted as a 'lack of blurriness' and has been examined in the context of visual saliency in other contexts. Researchers have documented depth-of-field effects, in which human attention is directed to the sharper areas of images, leaving the blurred parts as the background (Katz (1991), Marchesotti et al. (2015), Peterson et al. (2016), Loschky et al. (2014)). Several studies examine the blurriness and perception of depth (Held et al. (2010)) as well as the psychological sensitivity to image blur (Watson and Ahumada (2011)). In the human-computer interaction field, studies have focused on how image clarity can help users focus on relevant information (Enns and MacDonald (2013)). For instance, Veas et al. (2011) suggests that selected regions of image clarity can be used to guide viewers' attention. It stands to reason, therefore, that higher resolution images would attract investor attention more, and that campaigns with higher resolution images are more likely to be noticed. Further, the use of higher resolution images could signal stronger financial health of the company (better art-department, greater investment in graphic design). Either effect could potentially increase investor attention and funding.

3.2. Information Content/Reinforcement

While visual clarity measures address the visual quality of images, the second dimension we examine, information content/reinforcement, analyzes the actual information content depicted in the images. Our approach involves first determining whether images are informative, and then, for those that are, determining whether the identified information content is additive (incremental) to contemporaneously displayed text or merely reinforcing. This method, based on a machine learning classification algorithm, addresses a widely held concern regarding the inability of studies analyzing short term news or media effects on financial outcomes to control for potentially confounding information that is concurrently released.

3.2.1. Determining Whether the Information Content of an Image is Informative

To determine whether images contain informative content (constituting concrete identifiable objects) using an objective set of criteria, we use the the Google Vision API machine learning tool (Hereafter GV). To our knowledge, this is the first study to use this tool and to analyze the information content of images, particularly in these contexts. We classify images as 'informative' if GV identifies a concrete object and as 'uninformative' if it cannot.

For each image it analyzes, GV specifies a list of possible content descriptions or labels with corresponding estimated likelihoods (for those above 50%). For example, the tool can identify an 89% chance of an image representing Central Park, or a 98% chance that an image represents a pair of boots. The more 'factual' an image, the more descriptive in nature the type of information captured. For images that are more graphic in nature, such as simple textual logos, the GV algorithm returns labels which indicate that it cannot find a primary object.

——————Insert Figure 4 here——————

An example of an informative image is provided in Figure 4, which presents a stock photo of an elephant and a truncated list of the corresponding GV labels. The elephant (the information content of the image) is identified with high precision: the top label, 'Elephant and Mammoths', is attributed a 99% likelihood. Other high probability labels follow: Vertebrate (99%), Terrestial animal (98%) Wildlife (97%), Mammal (97%), Indian elephant (97%), and African elephant (96%). In contrast, we classify the the simple graphic in the figure (Meli Melo) as uninformative, since there is no concrete object identified, merely colors, texts and fonts. GV identifies this as a basic graphic image: 'font' (97%), 'text' (97%), 'logo' (86%), 'brown' (85%), 'brand' (69%), graphics (67%), and beige (55%). While GV is successful, its identifiers are merely descriptive and we classify the image and not informative (there is no concrete object in the graphic). This classification approach is flexible and easily customizable.

3.2.2. Determining the Information Reinforcement/Additivity of Images

For those images that are found to be informative, we now analyze the reinforcing nature of the images (to the information content of textual information). This classification is key in allowing us to distinguish between investor reactions to the different types of information and to disentangle the impact of imagery. Our reinforcement classification is based on a comparison between objectively derived image labels and words identified in accompanying text. Notably, our classification scheme does not include subjective qualification, nor does it attempt to evaluate the relevance or meaning of the informative content and potentially mitigates concerns regarding possibly confounding information that is concurrently released.

The methodology proposed is flexible in the degree to which textual keywords are identified, the number of reinforcement categories considered, and in the thresholds required for categorization. While in theory the information content of images can be either reinforcing, contradictory, additive or neutral, in this study we adopt a simple dichotomy and classify images as either additive or merely reinforcing. In its simplest and most objective form, the reinforcement criteria requires a single match between the image label and the textual information. In more complex situations, the image labels can be compared to text by keyword analysis or more advanced natural language processing techniques such as semantic similarity analysis or topic modeling. In our context, images are considered to be additive if there is no label-to textual word match and reinforcing otherwise.

Figure 5 presents two images, Kar-Go (Panel A) and 1854 Media (Panel B). We begin by determining that both images are informative by feeding the extracted images through GV and observing that the first label for each image is an object identified with high precision. Based on this classification (a necessary condition), we can now evaluate whether or the information content is additive or merely reinforcing to to the information content in the accompanying textual descriptions.

——————Insert Figure 5 here——————

Comparing the list of labels (partial list shown) for each image to the words in its corresponding textual description, we determine that Kar-Go is reinforcing (the overlapping word between the image label and the text in the Kar-Go pitch is 'vehicle') and that 1854 Media is additive (there are no overlapping words in the 1854 Media pitch.). We calculate visual clarity measures for both Kar-Go and 1854 Media and find that they are similar in singularity (4.6 and 4.5, respectively), but that they differ in their singularity score. For Kar-Go, singularity, but 1854 has a more negative singularity score than KarGo (-5.9 versus -5.0) reflecting the image's distinct focal points.

4. Data and Sample Description

Our hand-collected data comprise campaigns (pitches) from the two most active equity crowdfunding venues, Crowdcube and EquityNet, throughout our sample period.[12] The two platforms exhibit variations in size, funding information, and listing conventions. Notably, while funding information appears on all Crowdcube campaigns, it is available for only 73% of EquityNet campaigns (See Table 1). Visually, the two platforms also differ substantially, and examples of each are shown in Figure 6.

——————Insert Figure 6 here——————

Crowdcube displays large background cover images with inlaid logos and summary funding information, while images on EquityNet are restricted to small logos. Videos are allowed on both platforms, but are visually secondary, and when included, generally appear at the bottom of the page in smaller thumbnails. The distinct but relatively controlled visual environments encourage different visual attention tests.[13]

Table 1 (Panel A) describes the EquityNet sample construction process. We retrieve information on all 5,731 firm-generated campaigns (pitches) active as of September 07, 2020 and discard 1,568 campaigns that do not have both financial and firm level data. To control

[12]Companisto and Invesdor are also widely used but together comprise only 221 campaigns during our sample period and are therefore not included in our analysis.

[13]Firms generally list on one of these platforms and not another. Additionally, firms with several projects are rare in our sample.

for project newness and ensure that our data on funds raised is comparable across pitches, we discard 843 of the remaining 4,163 pitches that have been active for less than two years, resulting in a sample of 3,320 pitches.[14] We drop 1,745 pitches (53% of the remaining 3,320 pitches) that were unsuccessful in raising funds (zero dollars raised) after two years, resulting in a sample of 1,564 pitches. Since videos are generally not used (in over 75% of EquityNet pitches), we control for a potential video effect by discarding the 418 campaigns with videos.[15] From the resulting sample of 1,146, we identify 245 campaigns (21% of pitches) without perceptible logos (investors can see text only) and create a matched sample of 245 pitches with logos, resulting in a final sample of 490 EquityNet pitches.[16] For each campaign, we capture the campaign name, funding goal ($), funds raised ($), industry (business products, consumer products, financial services, industrial, IT& biotech, and media), number of employees, firm age, equity type (equity, convertible debt, debt financing, grant and loyalty) logo images, and the textual pitch description.

——————Insert Table 1 here——————

The average funding target for the 490 pitches in our final sample is roughly $1,930,000, the average amount raised is $674,000, and 92% of pitches are underfunded. Panel B of Table 1 provides a breakdown of characteristics for each of the two sub-samples (No-Logo and Logo) and reports a fairly balanced distribution across equity type, industry groups and funding status. Most of the projects (62%) seek equity financing, with convertible debt, debt, grant

[14]EquityNet campaigns tend to remain on the platform indefinitely, regardless of the degree to which they have achieved their funding goals. Since the platform does not provide information on campaign age or time listed, we collect the data in three snapshots over a 2 year window (September 15, 2018, May 22,2019, and September 07, 2020) to identify 3,320 campaigns that were active for at least two years. Less than two percent of campaigns in the merged data set changed logo images over this period. 78.5% of the 26,662 campaigns listed are categorized as community projects (not listed by the firms themselves) and therefore are not included in the sample.

[15]An alternative test was conducted using an unrestricted sample in which we control for the existence of videos, and the results are qualitatively similar. Further, since our focus of this study is on the effect of still images, including pitch pages with videos which may distract the investor's eye may obfuscate the resulting inferences. For this reason, the simple crowdfunding platforms utilized in this study are superior to other crowdfunding platforms such as in which videos are visually dominant and may also be embedded within text and images, thereby potentially obfuscating inferences derived from such tests.

[16]We ensure that the matched sample logos are clearly perceptible by conditioning on logo image size. We also create matched samples using propensity scores and the results are qualitatively similar.

and royalty financing following suit. Thirty-three percent of projects are self-reported by companies as Consumer Products, 17% as Financial Services, 13% as Business Products, 12% as Business Products, and the fall in the Media, Biotech, and IT industries.[17]

Table 2 describes the Crowdcube sample comprising the 976 campaigns (pitches) active as of July 27, 2020.[18] For each campaign, we capture the campaign name, funding goal (£), funds raised (£), firm age, logo images, background images, whether a pitch-related explanatory video is displayed, and the textual pitch description. All 976 pitches display logo images and videos, and 797 have background cover images that dominate the pitch visually. Fifty-four percent of pitches (526) have additional cash position information for two consecutive years. The average funding target (untabulated) for the entire sample is £416,490 and the average amount raised is £640,935. Unlike on EquityNet, the vast majority (88%) of projects are overfunded, with 2.8% (27) of projects exactly funded and 9.2% (91) underfunded. Projects are coded by type, and 355 projects are classified as Services, 197 as Food, 141 as Manufacturing, 123 as Entertainment, 81 as Finance, and 79 as Fashion. Untabulated results indicate that the amount contributed by each investor is typically small; while the mean project investment is £640,935, this figure is largely attributable to the the single largest investment in each campaign, averaging $93,000. The average investment amount for remaining investors is $1,620.

——————Insert Table 2 here——————

5. Empirical Analysis - The Equity Crowdfunding Application

We examine our visual attention methodology using three distinct testing arenas: EquityNet, Crowdcube, and a field experiment. Their suitability and limitations with regards

[17]EquityNet aggregates 39 self-reported industry sub-groupings into 12 larger industry groupings and we further aggregate these to 6 groupings to account for small group sizes. Results based on finer partitions yield qualitatively similar results.

[18]We are grateful to Crowdcube for granting us permission to download campaign information. Data collection was done in three rounds, September 15, 2018, May 22,2019, and July 27, 2020, and some cash variables were used from earlier rounds to verify 2020 data.

to testing different aspects of the methodology are summarized in Table 3. The EquityNet platform is most suitable to examine the image effect (whether or not the existence of a perceptible image affects financial decision making) using logos. Once we establish that equity crowdfunding campaigns with (perceptible) images indeed reach higher funding levels (after controlling for financial, sentiment and industry effects), we apply our visual attention methodology using data from Crowdcube's visually rich environment and directly test for the effects of clarity and information content/reinforcement on the degree to which investors are willing to commit their capital. Further, since on Crowdcube, investors see not only a logo, but also a main background (cover) image in addition to textual information, the image effect found in EquityNet can be disentangled from the logo effect. Finally, we validate our results with a field experiment in which we are able to control how information is viewed (only images, only text, or both). The field experiment allows us to observe directly as human participants make decisions about investment based upon their subjective perceptions of text, CrowdCube companies' background images, or a combination of the two. Doing this allows us to verify that the methodology we adopted for the first two portions of the study is accessing image variables that are relevant to human reporters as well as to machine algorithms. Further, it allows us to ask whether it is individuals' perceptions of these image variables that can affect investment decisions.

—————Insert Table 3 here—————

5.1. Testing Ground 1: EquityNet and The Image (Logo) Effect

EquityNet is identified as the most suitable platform to examine the image (logo) effect due to its controlled environment, in which only textual information and a logo (when available) are visible to investors. Since the No-Logo and Logo subsamples described in Section 4 do not include videos, we have a relatively controlled environment in which to determine whether an image effect exists (whether the existence of an image affects financial decision making). We estimate the effect of the existence of images (logos) on the dollar amount raised using

Equation (7):

$$Funds_i = \alpha + \beta_1 Logo_i + Controls_i + \Sigma EquityType + \Sigma Industry + \epsilon_i \tag{7}$$

where *Funds*, is the natural logarithm of the amount of funds raised, and *Logo*, is a dummy variable which equals one if the pitch has a logo and zero otherwise. We include several campaign-specific campaign variables. *Words*, the natural logarithm of the total number of words in the project description, proxies for the amount of textual information provided by insiders to prospective investors. *Tone* captures the sentiment of the textual campaign description and is calculated as the percentage difference in positive and negative words, as in Zhou et al. (2018).[19] *Funding Target* is the natural log of the campaign's funding target ($millions). Two other company specific variables are provided by EquityNet and are included as controls: *FirmAge*, the age of the company, and *Employees*, the number of employees at the firm. We also control for industry and equity type. Table 4 presents the regression variable descriptions and correlations.

——————Insert Table 4 here——————

Table 5 reports regression results and documents a strong image (logo) effect. *Logo* is positively and significantly associated with funds raised in all step-wise specifications. For the full specification (Column 5), *Logo* =0.65, and p <0.01. The economic significance of including a pitch logo on the platform this can be substantial; for pitches raising the average amount for the sample ($674,000), on average this translates to an increase of $438,100, and suggests a strong 'image effect'. A potential explanation for this phenomenon may be that these logos, as displayed on the main EquityNet pages, create a storefront window effect, in which store traffic is affected by the window displays, with investors clicking on to those pitches with logos.[20]

[19] The identification of positive and negative words is based on the Harvard-IV dictionary scheme. Replicating the analysis using Loughran and Mcdonald (2011) yields qualitatively similar results.

[20] See Oh and Petrie (2012) for a review of the literature examining the impact of store window displays on store entry decisions.We do not have access to click-through data to determine the number of clicks or

Campaign funding targets are also positively and significantly related to the amount of funds raised (*Funding Target*= 0.76;p <0.01), as is the number of employees at the firm (*Employees*= 0.40; p <0.01), consistent with our intuition, that larger or more established firms on average perform better in project financing. Equity type is also found to be significantly associated with the amount raised. Projects seeking debt and convertible debt financing are associated with higher funds raised than equity financing (the reference type), potentially capturing perceived risk of these projects.[21] Firm age, textual length, and sentiment are found not to be significant.

These results are consistent with our (limited) visual attention story, which we see as a natural extension of the attention (and limited attention) literature focusing on textual news and information. Hirshleifer and Teoh (2003), for example, conclude that investors are more likely to consider easy-processed information than more complicated items and that the form of presentation matters. Likewise, Knyazeva and Ivanov (2017) show that investors rely less on quantitative (financial) information than on other cues that attract their attention. In our context, since the market for equity crowdfunding is characterized by investors contributing relatively small amounts of money in exchange for a relatively small portion of the company's equity (Prive (2012), Wilmot (2018)), we argue that visual representation contributes to the (limited) attention of potential equity-crowdfunding campaign investors. Unlike venture capitalists, these investors are less likely to be sophisticated and possess less experience in valuing or assessing start-ups.

The inferences based on these results should be interpreted with some caution. First,

time spent on project landing page. Additionally, while EquityNet (and Crowdcube) could in theory be implementing black box internal ranking algorithms which would present "better" campaigns at more salient positions and therefore affect financing outcomes, this should not affect our results. Both platforms feature user-customizable drop down viewing controls, to order campaigns by zip code, project type, alpha ordering, funding goals, and more. We have verified that each investor is therefore presented with a different view and menu of campaigns, which invariably include both projects with and without high-resolution and colorful logos.

[21] These results are based on the sample constructed as described in Section 4, in which we remove roughly half of the sample which is unsuccessful in raising any money. Alternatively, we remove the restriction that pitches must have raised at least one dollar, and find qualitatively similar results.

endogeneity is often a concern. Second, documented associations could be confounded by potential omitted variables. The limited data availability on these private companies regarding project and entrepreneur quality impairs our ability to sufficiently control for project quality. On the other hand, the lack of widely available firm information (outside what is viewed on the EquityNet pages) can be helpful in providing a cleaner empirical setting within which to test our visual attention story. Third, the image effect we observe based on logos may be driven by logo-specific (and branding) factors. The use of different testing grounds in this study serves to partially mitigate some of these concerns. For example, the Crowdcube platform we examine next provides us with an opportunity to disentangle the image and logo effects.

5.2. Testing Ground 2: Crowdcube and Visual Attention

Unlike EquityNet, the image-intensive Crowdcube platform displays includes not only small logos and videos, but large and often colorful background images on crowdfunding campaign pages, allowing us to implement our visual attention methodology and explore how our two sets of criteria, visual clarity and information content/reinforcement, affect financial investment amounts. After revisiting the image effect, we apply our visual attention methodology by calculating clarity metrics for each background image and then identifying which images contain information content based on our image analysis algorithm. Finally, we determine whether the information content in each images is reinforcing or additive to the information content in accompanying textual information and examine investor reactions.

5.2.1. Crowdcube- Image versus Logo Effect

Results from the EquityNet sample revealed strong evidence of an image effect, with perceptible logos significantly associated with higher-funded campaigns. The Crowdcube platform, which includes both logos and colorful backgrounds, allows us to determine whether the EquityNet results were driven by an intrinsic ‘logo effect.’ Specifically, while all pitches require logos, background images are optional, and by testing whether or not the the existence of a background image is associated with investment allows us to disentangle, at least in part, the image effect documented for Equitynet from a potential logo effect, in which logos have intrinsic qualities that affect investor confidence or behavior. We estimate the background-driven image effect using Equation (8):

$$Funds_i = \alpha + \beta_1 LogoH_i + \beta_2 Background_i + Controls_i + \Sigma ProductGroup + \epsilon_i \tag{8}$$

Where *Funds* is the natural logarithm of the amount of funds raised, *Background* is a dummy variable which equals one if the campaign has a background image and zero if it does not. To proxy for the quality of logos, construct *LogoH*, which is a dummy variable equaling one if the logo image is in the top half of image quality (based on image size) and zero otherwise. We also control for financial variables constructed from information disclosed in the campaigns, including the campaign's target funding amount and two indicator variables: *CashInfo* equals one if cash position information exists for the company for the preceding two years and is zero otherwise, and *CashUp* equals one if there is an increase in cash position in the past two years, and is zero otherwise. These variables may be related to a company's financial stress, in that firms without cash positions displayed, and/or those for which cash positions do not increase may be perceived as those which are less likely to be able to realistically proceed with their projects. We also control for product type (Services, Food, Manufacturing, Entertainment, Finance, and Fashion). Table 6 presents descriptive statistics on the variables as well as correlations.

——————Insert Table 6 here——————

Table 7 presents the regression results. The image effect (based on the existence of Crowdcube background images) is positive and significant across all models. For example, in Model 4, *Background*= 0.13, p <0.01, translating to an average of £83,322 higher funding for pitches with background images (based on and average amount raised per pitch of £640,935). Finally, since firms with higher quality logos may be those that are more likely to have background images, we examine potential interactive effects between the two and present the results in

Models (5) and (6). The residual background effect remains significant.

——————Insert Table 7 here——————

Combined with the EquityNet results, we conclude that the existence of an image affects investor decision. On both platforms, we find that the image most prominently displayed to investors is associated with funds raised (on EquityNet this is the logo, and on Crowdcube this is the background cover image). Thus, while firm/project quality may be an omitted factor we cannot completely control for, the strong evidence of an image effect across these types of images indicates that the image effect does not appear to be an artifact of logo-specific attributes, is indicative of investor responsiveness to the existence of images in general, and may have policy impacts for firms seeking equity financing.

5.2.2. Crowdcube-Visual Attention Methodology- Clarity Metrics and Information Content

Having shown that the existence of backgrounds images is associated with larger Crowdcube investment amounts, we now turn our attention to the impact of visual clarity and information content measures on investors. We estimate the following equation, with two visual clarity measures, intensity and singularity:

$$Funds_i = \alpha + \beta_1 IBack_i + \beta_2 SBack_i + Controls_i + \Sigma ProductGroup + \epsilon_i \tag{9}$$

where *IBack*, and *SBack* are the background image intensity and singularity, respectively.[22]

Table 8 presents the results for the 797 Crowdcube campaigns that include background cover images. Both image clarity metrics are consistently significant across all models, with a positive (negative) association observed between intensity (singularity) and the amount of funds raised, suggesting that investors are more inclined to invest in projects with images that are more vibrant, and have more focal points. Economically, these numbers are also significant. For example, for the full specification, *IBack*=0.015 and *SBack*=-0.06 (Model (6)), and since

[22]Since Crowdcube uploads all backgrounds in a standard resolution format, that lack of variation in resolution precludes its measurement and inclusion in these tests.

the average Crowdcube campaign in our sample raises approximately (£640,935), each unit increase (decrease) in image intensity (singularity) is associated with on average, £103,831 (£37,302) larger project investment amounts.[23]

——————Insert Table 8 here——————

The funding target set by the crowdfunding campaign, *Funding Target*, is positive and significant across all models, indicating that larger projects are associated with higher investment levels in absolute terms. The sentiment variable, *Tone* does not load in any model, and while the variable *Words*, which captures the length of campaign descriptions, is positive and significant for Models (1), (3), and (5), significance disappears in all models that include cash position variables ((2), (4), and (6). This suggests that when cash variables are displayed on the pitch page, investor attention is less focused on text length. The combined results are consistent again with our visual attention story: investor attention seems to be drawn to the easiest-to-process information displayed. When large background images are available, such as on Crowdcube, information content in textual tone and length seem (at first blush) to be be less of a factor in determining investor behavior than that in the images themselves.

5.2.3. Crowdcube-Reinforcement/Additivity of Information Content

We further investigate our visual attention story by determining whether images with informationally additive content are associated with greater financial investment. We begin by differentiating between informative and uninformative images and then categorize informative images into those which are additive to the information content in accompanying textual information and those which are conversely, reinforcing in nature.[24] A total of 726 background images are classified as informative with the remainder deemed uninformative based on the

[23]This calculation is easily determined: (i.e., $ln(744,766) - ln(640,935)$ ∼0.15). (i.e., $ln(603,633) - ln(640,935)$ ∼-0.06).

[24]While the methodology proposed in this paper is useful in isolating the information content of pictorial representations from that in contemporaneously displayed text, we do not control for outside news releases, since we focus on the localized limited attention of equity crowdfunding investors in these private companies and since we are unable to ascertain when investors visit the site. Further, since crowdfunding project descriptions are relatively short, and the number of textual matches is relatively small, we limit our reinforcement categories to additive and reinforcing.

classification scheme outlined in the methodology section. Of those, 471 are found to be additive and 255 are reinforcing. For all 726 campaigns with informative background images, we now estimate the following equation:

$$Funds_i = \alpha + \beta_1 Clarity_i + \beta_2 Add_i + \beta_2 Add * Clarity_i + Controls_i + \Sigma ProductGroup + \epsilon_i \tag{10}$$

where *Clarity* is the continuous clarity score, calculated as the difference between the standardized values of intensity and singularity. [25] We estimate this score based on the results above showing the relative improvement based on the individual clarity components, singularity and intensity (higher levels on intensity and lower levels of singularity increase the overall clarity score). *Add* is a dummy variable that is 1 if the information in the background image is additive to the textual information and 0 if it is merely reinforcing, and *Add*Clarity* is an interactive term between *Add* and *Clarity*.[26]

————————Insert Table 9 here————————

Models (4), (5), an (6) include cash variables while (1), (2), and (3) do not. We find that *Clarity* has a positive and significant impact on funds for models (1) and (4). However, for specifications that include the additivity variable, we find that clarity is significant only when the background image is additive. Compared to the sample without additive backgrounds, a one unit increase in clarity increases funds by 18%.

Since additivity is a determining factor in our analysis, we now examine the effect of each of the clarity measures (intensity and singularity) for pitches with additive images. The results are reported in Table 10 and mirror those documented in Table 8. For the full specification (Model (2)), both clarity measures are significant at the one percent level, with *IBack*=0.22, and *SBack*=-0.08. When the additive measure is constructed using both the project summary

[25] We re-estimate the Equation (9) by using this clarity measure and find results are consistent with those reported by Table 8.

[26] We also conduct tests by interacting Add directly with IBack and SBack, and results are still consistent.

and the full available text (columns (3) and (4)), results are similar. [27]

We find that campaigns with additive backgrounds raise more funds on average than those with reinforcing images ($775,539 versus $633,764, respectively), and have lower funding targets ($436,291 versus $477,726). Notably, the average funding percentage for campaigns with additive images is significantly higher than for those with reinforcing text (untabulated). This result enforces our visual attention story- even when textual information is made available to investors along with the photos, campaigns with images that contain incremental information content tend to raise more funds. We conclude that the informational reinforcement dimension of visual attention is a critical factor in determining the value in imagery, and that images that are incrementally informative to accompanying textual information (additive images) are valued differently from those that are not (reinforcing images). Further, while we have earlier shown that the *existence* of images is valuable to financial investors, the visual characteristics of pictorial representations that affect investor choices differ across the two dichotomies: When images with informative content are also informationally reinforcing, campaigns tend to raise less funds on average, and visual clarity is not of import. The fact in that additive images, both singularity and intensity are correlated with financial success implies that when images bring something new to the table, they may receive more attention. And, in those cases, more vibrant images that have more focal points, may be those that help resolving uncertainty or conveying more information visually.

Taken together, the results in this paper are consistent with our simple (limited) visual attention story. First, we find strong evidence of an 'image effect' in which investor attention seems to be drawn to the easiest-to-process information displayed: when presented with a selection of crowdfunding campaigns that includes those with images and those without, investors will likely enter the 'stores' of those with perceptible images, which in turn tend to raise higher investment amounts. Second, in highly visual environments, when textual information

[27] Given the larger word count, and naturally larger match count, the reinforcement threshold is based on the median number of matches: informative images with above (below) median image label-to- textual word match are categorized as reinforcing (additive).

is also available, campaigns with images that provide incremental information content tend to raise higher funding percentages. Third, when images reinforce the information content in written text, visual clarity measures are not associated with higher investment amounts, but when the images bring something new to the table informationally, whether or not the images are vibrant and have several focal points appears to be important to investors, likely owing to the fact that such images may be more conducive to information extraction and uncertainty resolution.

5.3. Testing Ground 3: Field Experiment

The human participant study utilized CrowdCube data to explore non-investor perceptions and investment choices given differential information about companies. Participants were one hundred fifteen (115) college students enrolled in introductory psychology courses at a large (14,000 undergraduates) northeastern public university. All were college students between the ages of 18 and 22 years. As compensation for their participation participants received credit for completing a research requirement for their introductory courses.

All participants were informed of the purpose of the study verbally and in writing, and gave verbal and written consent for their participation. Each participant filled out one of three versions of an online questionnaire, taking 1.5-2 hours each. Questionnaires presented text only, background-picture only, or text plus picture data related to 35 CrowdCube companies.

Questionnaires asked participants whether they would invest in each company, why or why not, and asked for more specific information about student perceptions of the text and visual data with which they were presented. Salient to the current study, students were told that they had an imaginary $1000 to invest, and asked to allocate that money among the companies they had just learned about. For the questionnaires including background pictures and pictures plus textual descriptions, participants were asked to report what images they saw in the photo, and how the photo made them feel. They also reported their qualitative perceptions of the photographs, including information about image sharpness, and whether images were realistic or abstract. Critically, each participant defined image characteristics for themselves.

Because of this, the same background image might be rated subjectively as adding information to the textual description of a company or not. Therefore, rather than categorizing images as additive or not, data points from individual participants were considered as independent observations. Not every participant answered every question regarding image quality, leading to variation in degrees of freedom for individual analyses.

——————Insert Figure 7 here——————

Analysis of variance (ANOVA) revealed that human participants invested about the same amount of their imaginary money in companies regardless of whether they saw text only, background picture only, or text plus background picture. When a background picture was available, the factors that increased investment included additivity, singularity, intensity, and clarity of the images. Specifically, when the participant identified an image as providing information that went beyond the information in the text, they invested more money (additivity; $t1186 = 3.1$, $p \leq 0.01$). When the participant identified an image as having many (an average of four or more) versus few (average of 1) focal points (singularity; $t1290 = 2.1$, $p \leq 0.05$), they invested more money. When participants reported that an image was colorful versus grayscale (intensity; $t877 = 2.7$, $p \leq 0.01$), they invested nearly twice the amount of money. And when images were reported to have discernible objects (discernibility; $t1290 = 5$, $p \leq 0.001$), they invested more money. Realistic images attracted the same amount of investment as abstract images ($t1290 = 0.1$, NS)

——————Insert Figure 8 here——————

ANOVA reveals that singularity of focal point and image intensity interact, so that images that are either intense with multiple focal points or grayscale with singular focal points attract the most investment (no significant main effects; interaction F1,875= 9.1, $p \leq 0.0$). There are no significant interactions among other image variables.

——————Insert Figure 9 here——————

These data add to the information gleaned from previous analyses by revealing that discernible images (discernibility) evoke greater investment. Additionally, the interaction between focal point and intensity extend prior analysis indicating the importance of each of these stand-alone variables.

6. Conclusion

We argue that imagery is a missing factor in the current literature on attention and information analysis and highlight the importance of pictorial representation on financial decision making. Specifically, we explore whether incremental information conveyed by images contributes to the information set investors use. We introduce the concept of visual attention and propose a simple methodological framework within which to assess its impact in financial economics settings. Our approach in based on concrete (non-behavioral) metrics along two distinct sets of criteria: visual clarity and information content/reinforcement, which we believe jointly affect investors.

Our notion of visual attention can be seen as a natural extension of the attention (and limited attention) literature focusing on textual news and information. Hirshleifer and Teoh (2003), for example, conclude that investors are more likely to consider easy-processed information than more complicated items and that the form of presentation matters. A natural testing ground for our analytical framework therefore lies in the world of equity crowdfunding, which is characterized by typically smaller investors and visually rich settings. We show that both the level of clarity and the information content of images on the crowdfunding platform attract potential investors' attention at varying levels and affect the total amount raised, after controlling for both textual sentiment and financial information. The greater the information content in the image, and the more additive it is to the textual information, the larger the investment. The policy implications of our findings for firms are immediate: investing in image quality and information content can enhance financial performance in certain contexts. Beyond this, however, further research aimed at assessing whether these results hold for corporate products and issuances may be useful in guiding corporate policies.

Visual attention and imagery could potentially play an important role in the context of efficient markets. The Adaptive Expectations framework developed by Andy Lo (2017) asserts that the traditional efficient markets framework is flawed, in that it is incomplete, particularly in dynamic economies. Investors draw inferences using ‘human learning algorithms’ based on characteristics they are hardwired to deem important. Lo equates the ability to draw inferences to the degree of pixelization of images, which must be high enough to be able to discern specific patterns with which to make decisions. We posit that images add to the dimensionality of the set of characteristics people incorporate into their human learning algorithms, and that the pictorial representations can therefore bear on the efficiency of prices.

References

Abate, L. One Year of Equity Crowdfunding : Key Findings. Technical report, US Small Business Administration, Office of Advocacy, Economic Research Series, 2018.

Ackert, L. F., Church, B. K., and Deaves, R. Emotion and Financial Markets. Technical report.

Ahlers, G. K., Cumming, D., Günther, C., and Schweizer, D. Signaling in Equity Crowdfunding. *Entrepreneurship Theory and Practice*, 39(4):955–980, 7 2015.

Antonelli, D. Why Small Businesses Need Strong Logos, 2016. URL `https://www.entrepreneur.com/article/254602`.

Belfi, A. M., Vessel, E. A., Brielmann, A., Isik, A. I., Chatterjee, A., Leder, H., Pelli, D. G., and Starr, G. G. Dynamics of aesthetic experience are reflected in the default-mode network. *NeuroImage*, 188:584–597, 3 2019.

Ben-Rephael, A., Da, Z., and Israelsen, R. D. It depends on where you search: Institutional investor attention and underreaction to news. *Review of Financial Studies*, 2017.

Chemmanur, T. J. and Yan, A. Advertising, Attention, and Stock Returns. *SSRN Electronic Journal*, 2 2009.

Chen, Y. C., Colombatto, C., and Scholl, B. J. Looking into the future: An inward bias in aesthetic experience driven only by gaze cues. *Cognition*, 176:209–214, 7 2018.

Cowley, S. S.E.C. Gives Small Investors Access to Equity Crowdfunding - The New York Times, 2015. URL `https://www.nytimes.com/2015/10/31/business/dealbook/sec-gives-small-investors-access-to-equity-crowdfunding.html`.

Da, Z., Gurun, U. G., and Warachka, M. Frog in the pan: Continuous information and momentum. *Review of Financial Studies*, 2014.

Datta, R., Joshi, D., Li, J., and Wang, J. Z. Studying Aesthetics in Photographic Images Using a Computational Approach. Technical report, The Pennsylvania State University, 2006.

Dellavigna, S. and Pollet, J. M. Investor Inattention and Friday Earnings Announcements. *The Journal of Finance*, 64(2):709–749, 4 2009.

Desnoyer, M. and Wettergreen, D. Aesthetic Image Classification for Autonomous Agents. In *2010 20th International Conference on Pattern Recognition*, pages 3452–3455. IEEE, 8 2010.

Dougal, C., Engelberg, J., García, D., and Parsons, C. A. Journalists and the stock market, 3 2012.

Duarte, J., Siegel, S., and Young, L. Trust and credit: The role of appearance in peer-to-peer lending, 8 2012.

Enns, J. T. and MacDonald, S. C. The role of clarity and blur in guiding visual attention in photographs. *Journal of Experimental Psychology: Human Perception and Performance*, 39(2):568–578, 2013.

Estrin, S. and Khavul, S. Equity crowdfunding: a new model for financing entrepreneurship? *CentrePiece - The Magazine for Economic Performance*, 2016.

Fedyk, A. Front Page News: The Effect of News Positioning on Financial Markets for their invaluable guidance and advice. I am also grateful to. Technical report, Harvard University, 2018.

Focke, F., Ruenzi, S., Ungeheuer, M., Engelberg, J., Harvey, C., Odean, T., Mai, V., Schneemeier, J., and Schuster, P. Advertising, Attention, and Financial Markets a. Technical report.

Gentzkow, M., Kelly, B. T., and Taddy, M. Text As Data. 2017.

Giannini, R. C., Irvine, P. J., and Shu, T. Nonlocal Disadvantage: An Examination of Social Media Sentiment. 2011.

Google. Google Vision, 2019. URL `https://console.cloud.google.com/vision`.

Guzik, S. S. SEC Crowdfunding Rulemaking Under the Jobs Act – An Opportunity Lost? *SSRN Electronic Journal*, 3 2014.

Harris, D. B., Arnstine, D., O'leary, J. F., Kreitler, H., and Kreitler, S. The Psychology of the Arts. Technical Report 3, 1972.

Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., and Seidel, H. P. A statistical model of human pose and body shape. *Computer Graphics Forum*, 2009.

Held, R. T., Cooper, E. A., O'Brien, J. F., and Banks, M. S. Using blur to affect perceived distance and size. *ACM Transactions on Graphics*, 29(2):1–16, 3 2010.

Hirshleifer, D. and Teoh, S. H. Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics*, 36(1-3):337–386, 12 2003.

Hoberg, G. and Phillips, G. M. Text-Based Network Industries and Endogenous Product Differentiation. 2009.

Hsiao, S. H., Wang, Y. Y., Wang, T., and Kao, T. W. How social media shapes the fashion industry: The spillover effects between private labels and national brands. *Industrial Marketing Management*, 2019.

Hu, A. and Ma, S. Human Interactions and Financial Investment: A Video-Based Approach. *SSRN Electronic Journal*, 2020.

Huang, A. G., Tan, H., and Wermers, R. Institutional Trading around Corporate News: Evidence from Textual Analysis. 2014.

Jarvenpaa, S. L. Graphic Displays in Decision Making-The Visual Salience Effect. Technical report, 1990.

Katz, S. D. S. D. *Film directing shot by shot : visualizing from concept to screen.* Michael Wiese Productions in conjunction with Focal Press, 1991.

Knyazeva, A. and Ivanov, V. I. Soft and Hard Information and Signal Extraction in Securities Crowdfunding. *SSRN Electronic Journal*, 11 2017.

Larsen, V., Luna, D., and Peracchio, L. A. Points of View and Pieces of Time: A Taxonomy of Image Attributes. *Journal of Consumer Research*, 31(1):102–111, 6 2004.

Li, F. Annual Report Readability, Current Earnings, and Earnings Persistence *. *Journal of Accounting and Economics*, 45:221–247, 2008.

Liaukonyt, u., Zaldokas, A., Deuskar, P., George, L., Hirschey, N., Ibert, M., Li, W., Waldfogel, J., Weston, J., Gurun, U., Hong, H., Hwang, B.-H., Israelsen, R., Keloharju, M., Kučinskas, S., Lou, D., McGranaghan, M., Nagaraj, A., Peress, J., Schmidt, D., Shapiro, B., Yonker, S., and Wilbur, K. Background Noise? TV Advertising Affects Real Time Investor Behavior *. Technical report, 2018.

Liberti, J. M. and Petersen, M. A. Information: Hard and Soft*. *The Review of Corporate Finance Studies*, 8(1):1–41, 2019.

Loschky, L. C., Ringer, R. V., Johnson, A. P., Larson, A. M., Neider, M., and Kramer, A. F. Blur detection is unaffected by cognitive load. *Visual Cognition*, 22(3-4):522–547, 4 2014.

Lou, D. Attracting Investor Attention through Advertising. *Review of Financial Studies*, 27 (6):1797–1829, 6 2014.

Loughran, T. and Mcdonald, B. American Finance Association When Is a Liability Not a Liability? Textual Analysis. *Source: The Journal of Finance*, 66(1):35–65, 2011.

Loughran, T. and McDonald, B. Textual Analysis in Finance and Accounting: A Survey. *SSRN Electronic Journal*, 2014.

Luffarelli, J., Stamatogiannakis, A., and Yang, H. The Visual Asymmetry Effect: An Interplay of Logo Design and Brand Personality on Brand Equity. *Journal of Marketing Research*, 2019.

Madsen, J. and Niessner, M. Is Investor Attention for Sale? The Role of Advertising in Financial Markets. *Journal of Accounting Research*, 6 2019.

Malik, N., Vir Singh, P., Lee, D., and Srinivasan, K. When Does Beauty Pay. A Large Scale Image Based Appearance Analysis on Career Transitions. In *40th Annual ISMS Marketing Science Conference*, 2017.

Mamonov, S. and Malaga, R. Success factors in Title III equity crowdfunding in the United States. *Electronic Commerce Research and Applications*, 27:65–73, 1 2018.

Mandel, N. and Johnson, E. J. When Web pages influence choice: Effects of visual primes on experts and novice. Technical report, 2002.

Marchesotti, L., Murray, N., and Perronnin, F. Discovering Beautiful Attributes for Aesthetic Image Analysis. *International Journal of Computer Vision*, 113(3):246–266, 7 2015.

Mayer, E. J. Advertising, Investor Attention, and Stock Prices: Evidence from a Natural Experiment. Technical report, 2019.

Miller, E. and Kahn, B. Shades of Meaning: The Effect of Color and Flavor Names on Consumer Choice. *Journal of Consumer Research*, 32(1):86–92, 6 2005.

Mitchell, A. A. and Olson, J. C. Are Product Attribute Beliefs the Only Mediator of Advertising Effects on Brand Attitude? *Journal of Marketing Research*, 18(3):318, 8 1981.

Oh, H. and Petrie, J. How do storefront window displays influence entering decisions of clothing stores? *Journal of Retailing and Consumer Services*, 19(1):27–35, 1 2012.

Palmer, S., Schloss, K. B., and Sammartino, J. Visual Aesthetics and Human Preference. *Annual Review of Psychology*, 64, 2013.

Pech-Pacheco, J., Cristobal, G., Chamorro-Martinez, J., and Fernandez-Valdivia, J. Diatom autofocusing in brightfield microscopy: a comparative study. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, volume 3, pages 314–317. IEEE Comput. Soc, 2002.

Peterson, J., Ringer, R., Riter, M., Sisco, E., De La Torre, M., Subedi, S., and Loschky, L. The Effects of Blur/Clarity Contrast on Visual Selective Attention. *Journal of Vision*, 16 (12):994, 9 2016.

Plumlee, M., Brown, D., Hayes, R. M., and Marshall, R. S. Voluntary environmental disclosure quality and firm value: Further evidence. *Journal of Accounting and Public Policy*, 2015.

Pope, D. G. and Sydnor, J. R. What's in a Picture? Evidence of Discrimination from Prosper.com. *Source: The Journal of Human Resources*, 46(1):53–92, 2011.

Prive, T. Inside The JOBS Act: Equity Crowdfunding, 2012. URL `https://www.forbes.com/sites/tanyaprive/2012/11/06/inside-the-jobs-act-equity-crowdfunding-2/#1b451bac4b2e`.

Ravina, E. Love & Loans The Effect of Beauty and Personal Characteristics in Credit Markets. Technical report, 2019.

Scholars, J. W., Jack, J. ., and Beckwith, . ScholarlyCommons Predicting Success in Equity Crowdfunding. Technical report, University of Pennsylvania, 2016.

Schroeder, J. E. Visual Analysis of Images in Brand Culture. In Barbara J. Phillips, E. M., editor, *Go Figure: New Directions in Aadvertising Rhetoric*, pages 277–296. M.E. Sharpe, 2008.

Silvia, P. J. and Barona, C. M. Do People Prefer Curved Objects? Angularity, Expertise, and Aesthetic Preference. *Empirical Studies of the Arts*, 27(1):25–42, 1 2009.

Solomon, D. Selective Publicity and Stock Prices. *The Journal of Finance*, 67(2):599–638, 4 2012.

Tetlock, P. C. Does public financial news resolve asymmetric information? *Review of Financial Studies*, 23(9):3520–3557, 9 2010.

Veas, E. E., Mendez, E., Feiner, S. K., and Schmalstieg, D. Directing attention and influencing memory with visual saliency modulation. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, page 1471, New York, New York, USA, 2011. ACM Press.

Vessel, E. A., Maurer, N., Denker, A. H., and Starr, G. G. Stronger shared taste for natural aesthetic domains than for artifacts of human culture. *Cognition*, 179:121–131, 10 2018.

Vismara, S. Equity retention and social network theory in equity crowdfunding. *Small Business Economics*, 46(4):579–590, 4 2016.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., and von derHeydt, R. A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin*, 2012.

Walker, R. Why Fashion Brands All Use the Same-Style Font in Their Logos - Bloomberg, 2018. URL `https://www.bloomberg.com/news/articles/2018-11-20/why-fashion-brands-all-use-the-same-style-font-in-their-logos`.

Watson, A. B. and Ahumada, A. J. Blur clarified: A review and synthesis of blur discrimination. *Journal of Vision*, 11(5):10–10, 9 2011.

Wilmot, S. Slippery Way to Grab a Slice of the Next

Google - WSJ, 2018. URL https://www.wsj.com/articles/slippery-way-to-grab-a-slice-of-the-next-google-1525168533.

Yadav, A. Why Every Startup Should Host A Crowdfunding Campaign, 2018. URL https://www.forbes.com/sites/ajayyadav/2018/01/30/why-every-startup-should-host-a-crowdfunding-campaign/#7817863069ca.

Zhang, S., Lee, D. D., Singh, P. V., and Srinivasan, K. How Much is an Image Worth? Airbnb Property Demand Analytics Leveraging A Scalable Image Classification Algorithm. 2018. URL https://ssrn.com/abstract=2976021.

Zhou, M. J., Lu, B., Fan, W. P., and Wang, G. A. Project description and crowdfunding success: an exploratory study. *Information Systems Frontiers*, 20(2):259–274, 4 2018.

Tables and Figures

Table 1: **EquityNet Sample Description**
This table provides descriptive information on the EquityNet data used in this paper. Panel A tabulates the number of observations resulting from each step in the sample selection procedure. The final sample of 490 campaigns is formed by identifying the (245) pitches without perceptible logos and creating a matched sample of (245) perceptible pitches with logos. Panel B provides a descriptive breakdown of campaign (pitch) characteristics for the final sample of 490.

Panel A: Sample Selection

Procedure Description	Sample
Total Campaigns Collected	5,731
With Funding and Firm Information	4,163
Active for at Least Two Years	3,320
Have Raised any Funds	1,564
Does Not Display Video	1,146
Sub-sample of No logos	245
Matched sub-sample with logos	245
Final sample	**490**

Panel B: Final Sample Breakdown

	All	No-Logo	Logo
Sample	490	245	245
Equity Type			
# Equity	304	139	165
# Convertible Debt	66	28	38
# Debt	65	41	24
# Grant	39	27	12
# Royalty	16	10	6
Industry Group			
# Consumer Products	160	86	74
# Financial Services	85	59	26
# Business Products	65	31	34
# Industrial	58	29	29
# Media	47	19	28
# IT	75	21	54
Funding Status			
# Under Funded	452	227	225
# Exactly Funded	11	6	5
# Over Funded	27	12	15

Table 2: **Crowdcube Sample Description**
This table provides a descriptive breakdown of campaign (pitch) characteristics. # With Logo is the number of campaigns with logo images, # With Video is the number of campaigns with videos, # Overfunded is the number of campaigns for which the amount of funds raised exceeds the campaign financing target, and # Underfunded is the number of campaigns for which the amount of funds raised is less than the campaign financing target.

Data Breakdown	# Obs
Total Campaigns Collected	976
With Additional Financial Information	526
# With Logo	976
# With Video	997
# With Background Image	797
Funding Status	
# Under Funded	91
# Exactly Funded	27
# Over Funded	858
Product Group	
# Services	355
# Food	197
# Manufacturing	141
# Entertainment	123
# Finance	81
# Fashion	79

Table 3: **Testing Grounds for Visual Attention Measures and Methodology**
This table summarizes the suitability and limitations of the three testing grounds we use to test our visual attention measures and methodologies. Image Effect is broken down into either Logo effect or Background effect, Clarity includes both Intensity and Singularity, and Info includes the Informativeness measure as well as the ability to determine whether images are reinforcing or additive to the text information.

Testing Ground 1: EquityNet
Main Purpose: Test the Image Effect using Logos

Test:	Image Effect *Logo Effect*	Clarity	Info
Sample Suitability (Pros):			–
Clean visual layout and format.	✓	–	–
Images are limited to logos.	✓	X	X
Can compare pitches with and without logos.	✓	–	–
Most pitches do not have videos.	✓	–	–
Testing Limitations (Cons):			
Clarity measures can only be calculated for a subset of logos.			
Financial variables are extremely limited.			
Cannot disentangle Logo Effect from Image Effect.			
Small sample.			

Testing Ground 2: Crowdcube
Main Purpose: Test the Visual Attention Methodology

Test:	Image Effect *Background Effect*	Clarity	Info
Selection Criteria (Pros):			
Prominent Background Cover Image	✓	✓	✓
Other visual content is secondary; Can be Controlled.	✓	✓	✓
Can disentangle Logo Effect from Image Effect.	✓		
Testing Limitations (Cons):			
Small sample.			

Testing Ground 3: Surveys
Main Purpose: Validate Results

Test:	Image Effect *Background Effect*	Clarity	Info
Selection Criteria (Pros):			
Controlled Environment	✓	✓	✓
Prominent Background Cover Image	✓	✓	✓
Flexible Research Design (display only cover; only text; both)	✓	✓	✓
Testing Limitations (Cons):			
Experimental Study.			

Table 4: **EquityNet Variable Descriptive Statistics and Correlations**
This table presents summary statistics and correlations for the regression variables used in the EquityNet sample. Panel A presents summary statistics and Panel B presents correlations. The variable *Funds* is the natural logarithm of the dollar amount of funds raised. *Logo* is a dummy variable which equals one when the logo image file size is in the top image file size quintile and zero when it is in the bottom quintile, *Video* is an indicator for whether the campaign has a video or not, and equals one when there is a video and zero otherwise, *Words* is the natural logarithm of the total number of words in the project description, *Tone* is calculated as the positive word count minus the negative word count divided by the total count of positive and negative words, *FirmAge* is the natural logarithm of the age of the company, in years, *Funding Target* is the natural log of the total funding target ($ million), and *# Employees* is the natural logarithm of the company's number of employees. Pearson correlations are given in the lower half triangle. Spearman correlations are given in the upper half triangle. * denotes significance at 5%.

Panel A. EquityNet Variable Descriptions

Variables	**# of Obs.**	**Mean**	**Std Dev**	**Minimum**	**25th Perc**	**Median**	**75th Perc**	**Maximum**
Funding Target ($mil)	490	1.93	5.52	0.00	0.20	0.50	1.20	60.00
Funding Target	490	13.30	1.30	11.51	12.43	13.12	14.00	17.05
Funds	490	10.95	2.24	4.30	9.62	11.00	12.43	16.14
Words	490	4.80	0.76	2.77	4.32	4.87	5.44	5.93
Tone	490	0.24	0.15	−0.07	0.15	0.23	0.33	0.70
FirmAge	490	1.93	0.50	0.69	1.61	1.95	2.08	3.47
# Employees	490	1.46	0.79	0.00	1.10	1.39	1.79	3.86

Panel B. EquityNet Variable Correlations

Funds	*Funds*	*Logo*	*Words*	*Tone*	*FirmAge*	*# Employees*	*Funding Target*
Funds	1	0.27***	0.13**	0.07	0.17***	0.30***	0.58***
Logo	0.27***	1	0.25***	0.00	0.11*	0.15***	0.17***
Words	0.12**	0.27***	1	0.12**	0.00	0.07	0.12**
Tone	0.02	−0.01	0.09*	1	−0.01	0.12**	0.07
FirmAge	0.14**	0.06	−0.04	−0.02	1	0.24***	0.11*
# Employees	0.29***	0.14**	0.04	0.05	0.22***	1	0.20***
Funding Target	0.55***	0.17***	0.13**	0.05	0.09*	0.23***	1

Table 5: **EquityNet Image Effect Regression Results** This table presents regression results measuring whether the presence of EquityNet images (logos) affects the total funds raised. The reported parameters are coefficients from the following regression: $Funds_i = \alpha + \beta_1 Logo_i + Controls_i + \Sigma EquityType + \Sigma Industry + \epsilon_i$. The five models correspond to step-wise specifications. The dependent variable *Funds* is the natural logarithm of the dollar amount of funds raised. *Logo* is a dummy variable equal to one if the pitch has a logo and 0 otherwise, *Funding Target* is the natural log of the total funding target ($ million), *Words* is the natural logarithm of the total number of words in the project description, *Tone* is calculated as the positive word count minus the negative word count divided by the total count of positive and negative words, and *Industry Group* includes eight different industry groupings. All models are estimated by using the ordinary least squares (OLS) model with robust standard errors. Standard-errors are in square brackets, and significance levels are indicated as follows: $^{*}-$ 10%, $^{**}-$ 5%, $^{***}-$ 1%.

	Funds				
	(1)	(2)	(3)	(4)	(5)
Intercept	−1.23 **	−1.87 ***	−1.54 **	−1.19	0.58
	[0.49]	[0.67]	[0.72]	[0.90]	[0.83]
Logo	0.83***	0.80***	0.74***	0.67***	0.65***
	[0.13]	[0.15]	[0.13]	[0.05]	[0.05]
Funding Target	0.88***	0.87***	0.82***	0.79***	0.76***
	[0.04]	[0.04]	[0.05]	[0.02]	[0.04]
Words	0.01	0.02	0.03	0.05	0.05
	[0.15]	[0.17]	[0.15]	[0.14]	[0.15]
Tone	0.00	0.02	−0.07	0.05	0.18
	[0.17]	[0.17]	[0.29]	[0.22]	[0.24]
FirmAge		0.38*	0.26	0.28	0.23
		[0.22]	[0.19]	[0.18]	[0.18]
# Employees			0.41***	0.39***	0.40***
			[0.09]	[0.10]	[0.10]
ConvDebt				0.19***	0.19***
				[0.01]	[0.02]
Debt				0.30***	0.29***
				[0.02]	[0.02]
Grant				−1.40 ***	−1.41 ***
				[0.02]	[0.02]
Royalty				−0.40 ***	−0.43 ***
				[0.05]	[0.04]
Industry Group					included
Observations	490	490	490	490	490
Adjusted R^2	0.33	0.33	0.35	0.38	0.38

Table 6: **Crowdcube Variable Descriptive Statistics and Correlations**
This table presents summary statistics and correlations for the regression variables used in the Crowdcube sample. Panel A presents summary statistics and Panel B presents correlations. The variable *Funds* is the natural logarithm of the dollar amount of funds raised. *Words* is the natural logarithm of the total number of words in the project description, *Tone* is calculated as the positive word count minus the negative word count divided by the total count of positive and negative words, *FirmAge* is the natural logarithm of the age of the company, in years, *Funding Target* is the natural log of the total funding target (£million), *IBack* is the natural logarithm of background image intensity, *SBack* is the natural logarithm of the background focus score times (-1), *CashUp* is an indicator for whether the campaign displays an increase in cash position in the past two years, and equals one when it does and zero otherwise, and *Reinforcing* is an indicator which equals one if the informative content of the background image is reinforcing of the text and zero otherwise. Pearson correlations are given in the lower half triangle. Spearman correlations are given in the upper half triangle. * denotes significance at 5%.

Panel A. Crowdcube Variable Descriptions

Variables	**# of Obs.**	**Mean**	**Std Dev**	**Minimum**	**25th Perc**	**Median**	**75th Perc**	**Maximum**
Funds	976	2.74	1.09	0.15	1.93	2.62	3.48	5.62
IBack	797	4.86	0.28	3.93	4.72	4.93	5.08	5.22
SBack	797	−5.81	1.03	−7.73	−6.55	−5.95	−5.21	−2.73
Words	976	3.89	0.29	2.89	3.76	3.87	4.04	6.82
Tone	976	0.26	0.17	−0.08	0.15	0.24	0.37	0.74
Reinforcing	726	0.35	0.48	0.00	0.00	0.00	1.00	1.00
FirmAge	976	1.36	0.60	0.00	1.10	1.39	1.79	2.66
CashInfo	526	0.21	0.41	0.00	0.00	0.00	0.00	1.00
CashUp	526	0.35	0.48	0.00	0.00	0.00	1.00	1.00
Funding Target (£)	976	0.29	0.28	0.00	0.10	0.18	0.41	3.14

Panel B. Crowdcube Variable Correlations

Variables	*Funds*	*IBack*	*SBack*	*Words*	*Tone*	*Reinforcing*	*FirmAge*	*CashInfo*	*CashUp*	*Funding Target*
Funds	1	0.09	−0.06	0.06	−0.03	−0.02	0.41	−0.11	0.10	0.92
IBack	0.07	1	−0.11	0.01	0.05	0.05	−0.01	0.06	−0.02	0.09
SBack	−0.07	−0.19	1	−0.14	−0.05	−0.06	−0.04	−0.02	0.01	−0.05
Words	0.04	0.03	−0.12	1	0.02	0.06	−0.06	−0.05	−0.05	0.06
Tone	−0.05	0.05	−0.05	0.07	1	0.04	0.01	−0.01	0.01	−0.01
Reinforcing	−0.02	0.03	−0.06	0.07	0.03	1	0.04	−0.06	0.15	−0.02
FirmAge	0.40	−0.01	−0.04	−0.02	0.01	0.04	1	−0.22	0.09	0.36
CashInfo	−0.11	0.07	−0.02	−0.03	−0.01	−0.06	−0.23	1	−0.38	−0.13
CashUp	0.10	−0.02	0.00	−0.01	0.01	0.15	0.10	−0.38	1	0.12
Funding Target	0.80	0.04	−0.03	0.03	−0.05	−0.02	0.31	−0.10	0.08	1

Table 7: **Crowdcube Image Effect Regression Results** This table presents regression results measuring whether the presence of images (Crowdcube background covers) affects the total funds raised. The reported parameters are the intercept and coefficients from the following regression: $Funds_i = \alpha + \beta_1 LogoH_i + \beta_2 Background_i + Controls_i + \Sigma ProductGroup + \epsilon_i$. The four models correspond to step-wise specifications. The dependent variable *Funds* is the natural logarithm of the amount of funds raised (£). *Background* is a dummy variable which equals one if the campaign has a background image and zero if it does not, *LogoH* is a dummy variable which equals one when the logo image file size is in the top half of image file size and zero when it is in the bottom half, *Words* is the natural logarithm of the total number of words in the project description, *Tone* is calculated as the positive word count minus the negative word count divided by the total count of positive and negative words, *Funding Target* is the natural log of the total funding target (£million), *FirmAge* is the natural logarithm of the age of the company, in years, *CashInfo* is an indicator for whether the campaign displays cash info in either of the past two years, and equals one when it does and zero otherwise, *CashUp* is an indicator for whether the campaign displays an increase in cash position in the past two years, and equals one when it does and zero otherwise, *Background*LogoH* is an interactive term between *Background* and *LogoH*, and *Product Group* includes the six product groups: Services, Food, Manufacturing, Entertainment, Finance, and Fashion. All models are estimated by using the ordinary least squares (OLS) model with robust standard errors. Standard-errors are in square brackets, and significance levels are indicated as follows: $^{*}-$ 10%, $^{**}-$ 5%, $^{***}-$ 1%.

	Funds					
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.93***	1.21***	0.94***	1.22***	0.94***	1.21***
	[0.13]	[0.09]	[0.12]	[0.09]	[0.31]	[0.29]
Background	0.26***	0.14***	0.23***	0.13***	0.22***	0.14**
	[0.01]	[0.04]	[0.01]	[0.04]	[0.07]	[0.07]
LogoH			0.10***	0.08***	0.13	0.10
			[0.01]	[0.02]	[0.11]	[0.10]
Words	0.15***	0.03	0.14***	0.02	0.14*	0.02
	[0.01]	[0.03]	[0.01]	[0.03]	[0.07]	[0.07]
Tone	−0.09 ***	−0.13 *	−0.08 ***	−0.13 **	−0.08	−0.13
	[0.02]	[0.07]	[0.01]	[0.06]	[0.13]	[0.14]
Funding Target	2.99***	3.55***	3.00***	3.54***	3.00***	3.54***
	[0.11]	[0.24]	[0.11]	[0.24]	[0.08]	[0.11]
FirmAge		0.19***		0.20***		0.20***
		[0.06]		[0.06]		[0.04]
CashInfo		0.03		0.03		0.03
		[0.12]		[0.11]		[0.06]
CashUp		0.05**		0.05**		0.05
		[0.02]		[0.02]		[0.05]
*Background*LogoH*					−0.03	−0.03
					[0.12]	[0.12]
Product Group	included	included	included	included	included	included
Observations	976	526	976	526	976	526
Adjusted R^2	0.65	0.74	0.65	0.74	0.65	0.74

Table 8: **The Effects of Image Clarity (Intensity and Singularity) on Financial Success** This table presents regression results measuring whether Crowdcube background image clarity metrics (singulairty and intensity) affect the total funds raised in crowdfunding campaigns. The sample includes those 797 campaigns for which background images are displayed. The reported parameters are the intercept and coefficients from the following regression: $Funds_i = \alpha + \beta_1 IBack_i + \beta_2 SBack_i + Controls_i + \Sigma ProductGroup + \epsilon_i$. The six models correspond to step-wise specifications. The dependent variable *Funds* is the natural logarithm of the amount of funds raised (£). *IBack* is the natural logarithm of background image intensity, *SBack* is the natural logarithm of the background focus score times (-1), *Words* is the natural logarithm of the total number of words in the project description, *Tone* is calculated as the positive word count minus the negative word count divided by the total count of positive and negative words, *Funding Target* is the natural log of the total funding target (£million), *FirmAge* is the natural logarithm of the age of the company, in years, *CashInfo* is an indicator for whether the campaign displays cash info in either of the past two years, and equals one when it does and zero otherwise, *CashUp* is an indicator for whether the campaign displays an increase in cash position in the past two years, and equals one when it does and zero otherwise, and *Product Group* includes the six product groups: Services, Food, Manufacturing, Entertainment, Finance, and Fashion. All models are estimated by using the ordinary least squares (OLS) model with robust standard errors. Standard-errors are in square brackets, and significance levels are indicated as follows: $^{*}-$ 10%, $^{**}-$ 5%, $^{***}-$ 1%.

	Funds					
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	−0.09	0.44	0.46	1.10	−0.10	0.42
	[1.03]	[1.23]	[0.71]	[1.20]	[0.94]	[1.15]
IBack	0.15**	0.20***			0.12**	0.15***
	[0.07]	[0.01]			[0.05]	[0.01]
SBack			−0.05 **	−0.07 ***	−0.04 ***	−0.06 ***
			[0.02]	[0.003]	[0.02]	[0.004]
Words	0.21*	0.001	0.18**	−0.03	0.19**	−0.03
	[0.11]	[0.26]	[0.08]	[0.25]	[0.09]	[0.25]
Tone	−0.12	−0.08	−0.12	−0.07	−0.13	−0.08
	[0.10]	[0.12]	[0.10]	[0.16]	[0.10]	[0.14]
Funding Target	2.68***	3.70***	2.69***	3.71***	2.68***	3.70***
	[0.41]	[0.71]	[0.41]	[0.70]	[0.40]	[0.70]
FirmAge	0.34***	0.23***	0.33***	0.23***	0.33***	0.23***
	[0.01]	[0.03]	[0.02]	[0.02]	[0.01]	[0.03]
CashInfo		0.09**		0.10**		0.09*
		[0.05]		[0.04]		[0.05]
CashUp		0.06		0.06		0.06
		[0.08]		[0.08]		[0.08]
Product Group	included	included	included	included	included	included
Observations	797	389	797	389	797	389
Adjusted R^2	0.65	0.74	0.65	0.74	0.65	0.74

Table 9: **The Effects of Image Additivity on Financial Success** This table reports the effect of the additivity of images on the total funds raised for the sample of 726 campaigns with informative background images. The first three columns present results of the model estimated for the full sample, and the last three columns are estimated for the set of pitches with cash information. *Funds* is the natural logarithm of the amount of funds raised (£). *Clarity* is the continuous clarity score, calculated as the difference between the standardized values of intensity and singularity. *Add* is a dummy variable that is 1 if the information in the background image is additive to the textual information and 0 if it is merely reinforcing. *Add*Clarity* is an interactive term between *Additivity* and *Clarity*. *IBack* is the natural logarithm of background image intensity, *SBack* is the natural logarithm of the background focus score times (-1), *Words* is the natural logarithm of the total number of words in the project description, *Tone* is calculated as the positive word count minus the negative word count divided by the total count of positive and negative words, *Funding Target* is the natural log of the total funding target (£million), *FirmAge* is the natural logarithm of the age of the company, in years, *CashInfo* is an indicator for whether the campaign displays cash info in either of the past two years, and equals one when it does and zero otherwise, *CashUp* is an indicator for whether the campaign displays an increase in cash position in the past two years, and equals one when it does and zero otherwise, and *Product Group* includes the six product groups: Services, Food, Manufacturing, Entertainment, Finance, and Fashion. All models are estimated by using the ordinary least squares (OLS) model with robust standard errors. Standard-errors are in square brackets, and significance levels are indicated as follows: $^{*}-$ 10%, $^{**}-$ 5%, $^{***}-$ 1%.

	Funds					
	Without Cash Information			With Cash Information		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.79	0.63	0.72	1.48	1.44	1.53
	[0.56]	[0.71]	[0.60]	[1.21]	[1.37]	[1.24]
Clarity	0.16**		0.05	0.18***		0.03
	[0.07]		[0.13]	[0.01]		[0.11]
Add		0.04	0.05		−0.05	−0.04
		[0.06]	[0.07]		[0.07]	[0.07]
*Add*Clarity*			0.18**			0.26*
			[0.07]			[0.14]
Words	0.18**	0.22**	0.19**	−0.02	−0.001	−0.02
	[0.08]	[0.10]	[0.08]	[0.25]	[0.27]	[0.25]
Tone	−0.14 *	−0.12	−0.14 *	−0.09	−0.08	−0.11
	[0.09]	[0.10]	[0.08]	[0.12]	[0.12]	[0.11]
Funding Target	2.67***	2.68***	2.66***	3.69***	3.71***	3.69***
	[0.43]	[0.44]	[0.43]	[0.71]	[0.72]	[0.69]
FirmAge	0.32***	0.32***	0.32***	0.23***	0.23***	0.22***
	[0.004]	[0.01]	[0.01]	[0.03]	[0.02]	[0.02]
CashInfo				0.09*	0.10***	0.09
				[0.05]	[0.04]	[0.06]
CashUp				0.06	0.05	0.05
				[0.08]	[0.09]	[0.09]
Product Group	included	included	included	included	included	included
Observations	726	726	726	388	388	388
Adjusted R^2	0.65	0.64	0.65	0.74	0.73	0.74

Table 10: **Informationally Additive Pitches (Short and Expanded Textual Information)** This table presents regression results measuring whether Crowdcube background image clarity metrics (intensity and singularity) affect the total funds raised in crowdfunding campaigns, for the 471 campaigns with additive background images. Results are reported separately for two sets of analyses which differ based on which textual descriptions we compare the image labels to (in determining the additivity of image information content). The first two columns contain results for regressions conducted on a sample of campaigns for which the reinforcement classification was based on project summary descriptions and the last two columns contain results for regressions conducted on a sample of campaigns for which the reinforcement classification was based on full project descriptions. *IBack* is the natural logarithm of background image intensity, *SBack* is the natural logarithm of the background focus score times (-1), *Words* is the natural logarithm of the total number of words in the project description, *Tone* is calculated as the positive word count minus the negative word count divided by the total count of positive and negative words, *Funding Target* is the natural log of the total funding target (£million), *FirmAge* is the natural logarithm of the age of the company, in years, *CashInfo* is an indicator for whether the campaign displays cash info in either of the past two years, and equals one when it does and zero otherwise, *CashUp* is an indicator for whether the campaign displays an increase in cash position in the past two years, and equals one when it does and zero otherwise, and *Product Group* includes the six product groups: Services, Food, Manufacturing, Entertainment, Finance, and Fashion. All models are estimated by using the ordinary least squares (OLS) model with robust standard errors. Standard-errors are in square brackets, and significance levels are indicated as follows: $^{*}-$ 10%, $^{**}-$ 5%, $^{***}-$ 1%.

	Funds			
	Project Summary Text		Full Project Description Text	
	(1)	(2)	(3)	(4)
Intercept	−1.61	−1.23	−1.28	0.77
	[1.21]	[1.96]	[0.88]	[0.35]
IBack	0.25***	0.22***	0.28***	0.17***
	[0.06]	[0.05]	[0.06]	[0.05]
SBack	−0.04 ***	−0.08 ***	−0.11 **	−0.07 ***
	[0.02]	[0.02]	[0.05]	[0.01]
Words	0.39***	0.18	0.40***	−0.12 *
	[0.14]	[0.41]	[0.15]	[0.07]
Tone	−0.12 ***	−0.19	−0.11 ***	−0.03
	[0.03]	[0.17]	[0.04]	[0.18]
Funding Target	3.24***	3.48***	3.25***	3.52***
	[0.11]	[0.58]	[0.11]	[0.60]
FirmAge	0.24**	0.32***	0.25**	0.27***
	[0.11]	[0.005]	[0.11]	[0.01]
CashInfo		0.17		0.18
		[0.21]		[0.13]
CashUp		0.08		0.08
		[0.14]		[0.07]
Product Group	included	included	included	included
Observations	471	221	363	194
Adjusted R^2	0.73	0.72	0.71	0.70

Figure 1: **Methodology** This figure presents the simple methodology used to analyze the impact of images

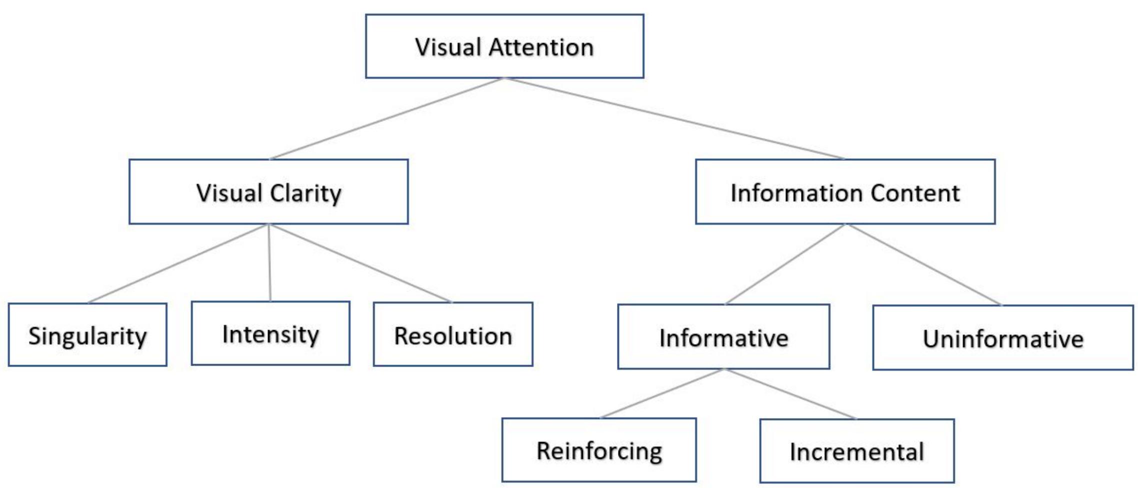

Figure 2: **Image Intensity** This figure presents two crowdfunding background images. The image in Panel A has a higher intensity score than the one in Panel B.

(a) Panel A

(b) Panel B

Source: Images from crowdcube.com: Panel A: Doorsteps.co.uk; Panel B: Redemption Brewing Company.

Figure 3: **Image Singularity** This figure presents two crowdfunding background images in our sample. The image in Panel A has a higher singularity score than the one in Panel B.

(a) Panel A

(b) Panel B

Source: Images from crowdcube.com: Panel A: Nibbling Jewellery; Panel B: Playbrew.

Figure 4: **Informative and Non Informative Images** This figure presents two images and corresponding Google Vision algorithm output. The image in Panel A is classified as informative and the image in Panel B is not.

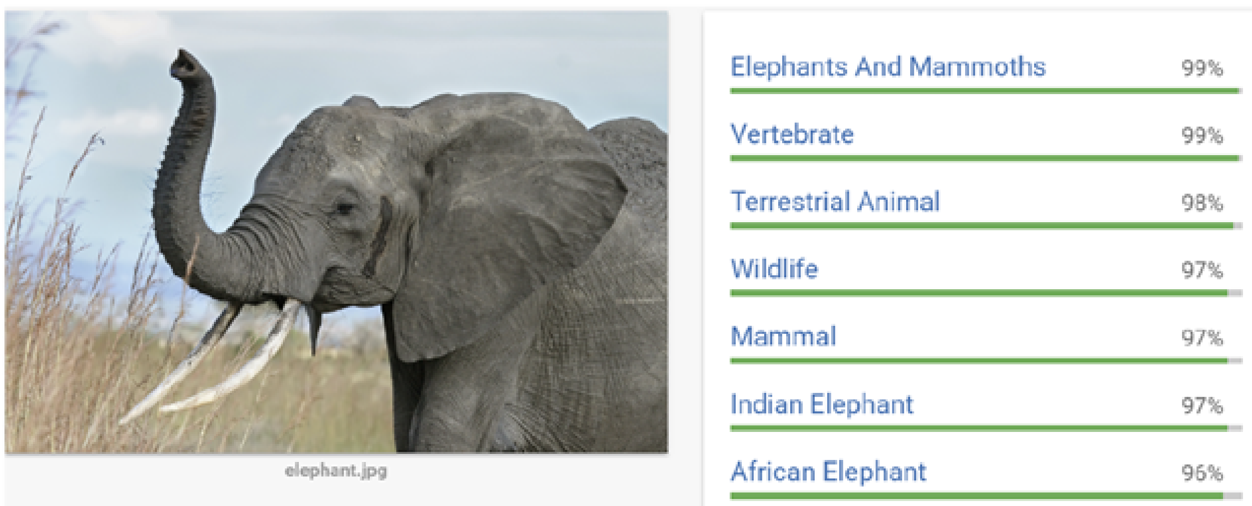

(a) Panel A

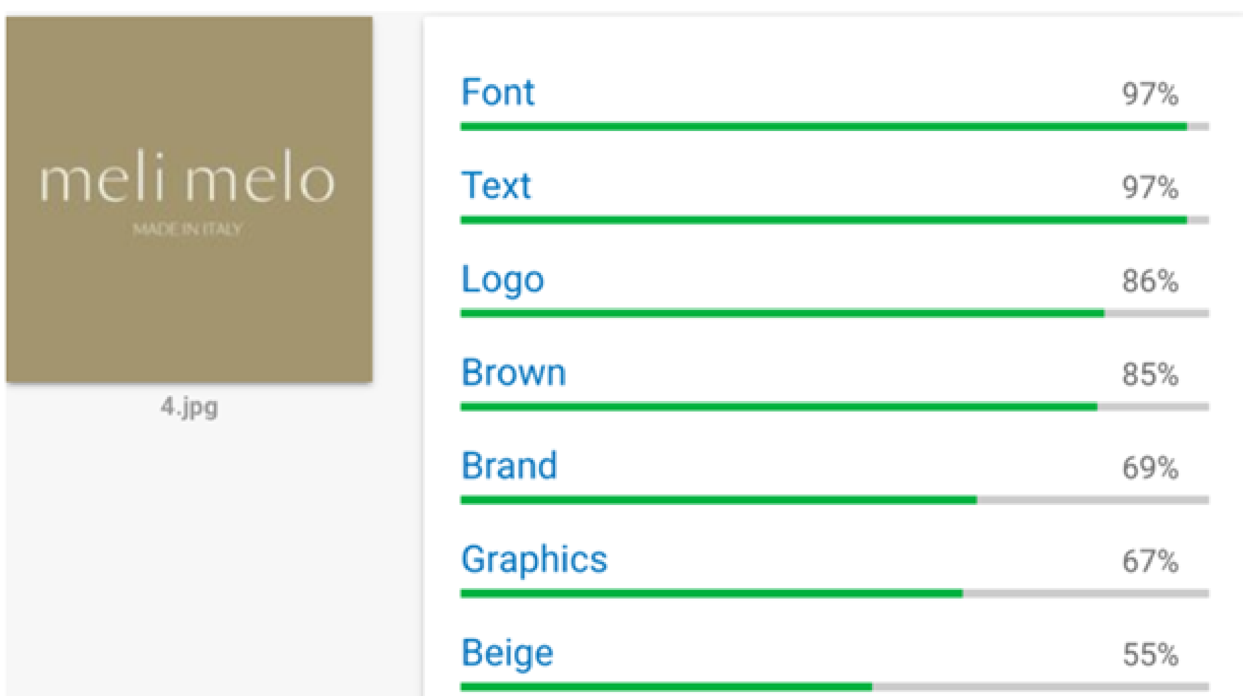

(b) Panel B

Source: Image from Wikipedia; output from Google Vision

Figure 5: **Information Additivity and Reinforcement** This figure presents two images from equity crowdfunding venues. The image in Panel A is classified as 'Reinforcing' and the image in Panel B is classified as 'Additive' (to the information provided textually).

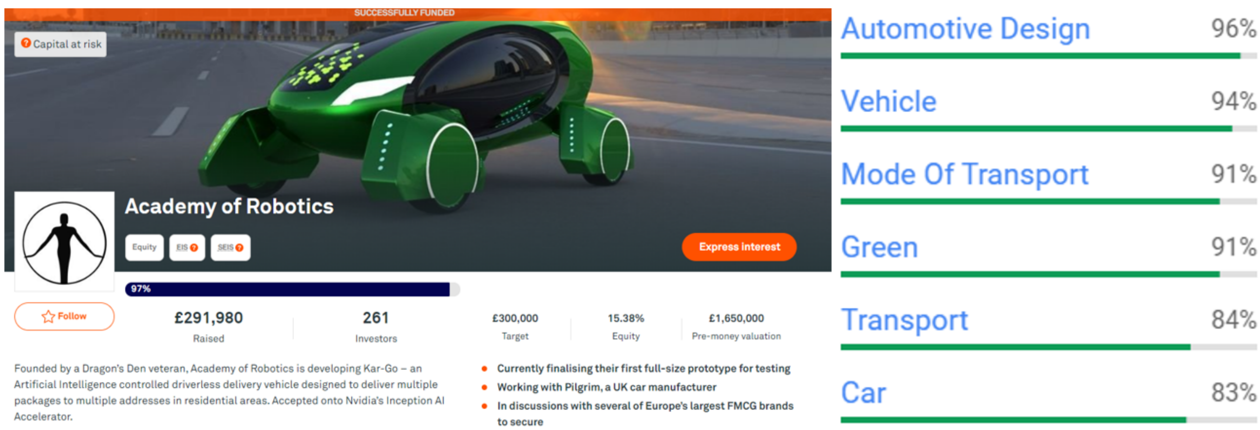

(a) Panel A

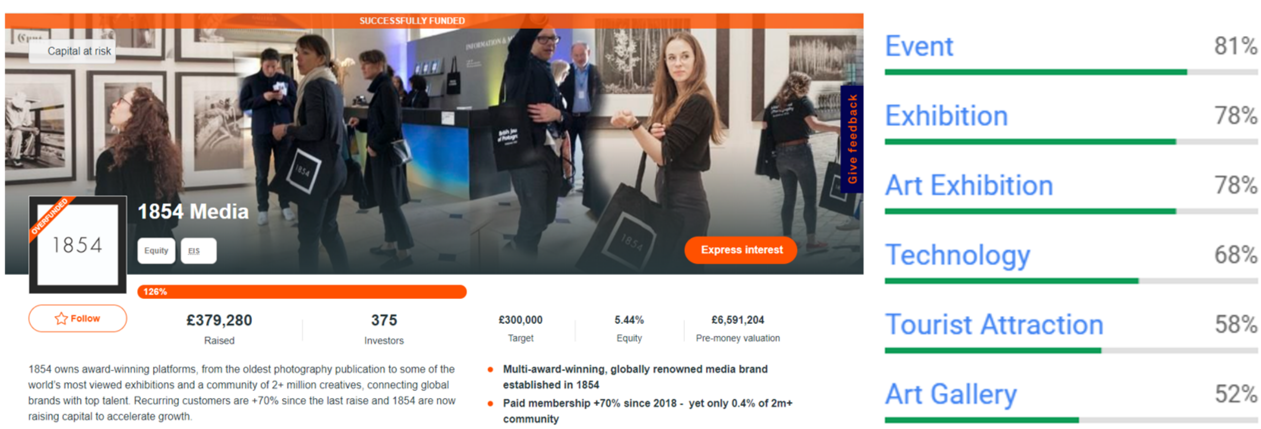

(b) Panel B

Source: Images from crowdcube.com: Panel A: Academy of Robotics (Kar-Go); Panel B: 1854 Media Ltd.

Figure 6: **Platform Visual Environments** This figure highlights the different visual environments of the EquityNet and Crowdcube Platforms. Panel A presents a representative EquityNet crowdfunding campaign page and Panel B presents a representative Crowdcube campaign page.

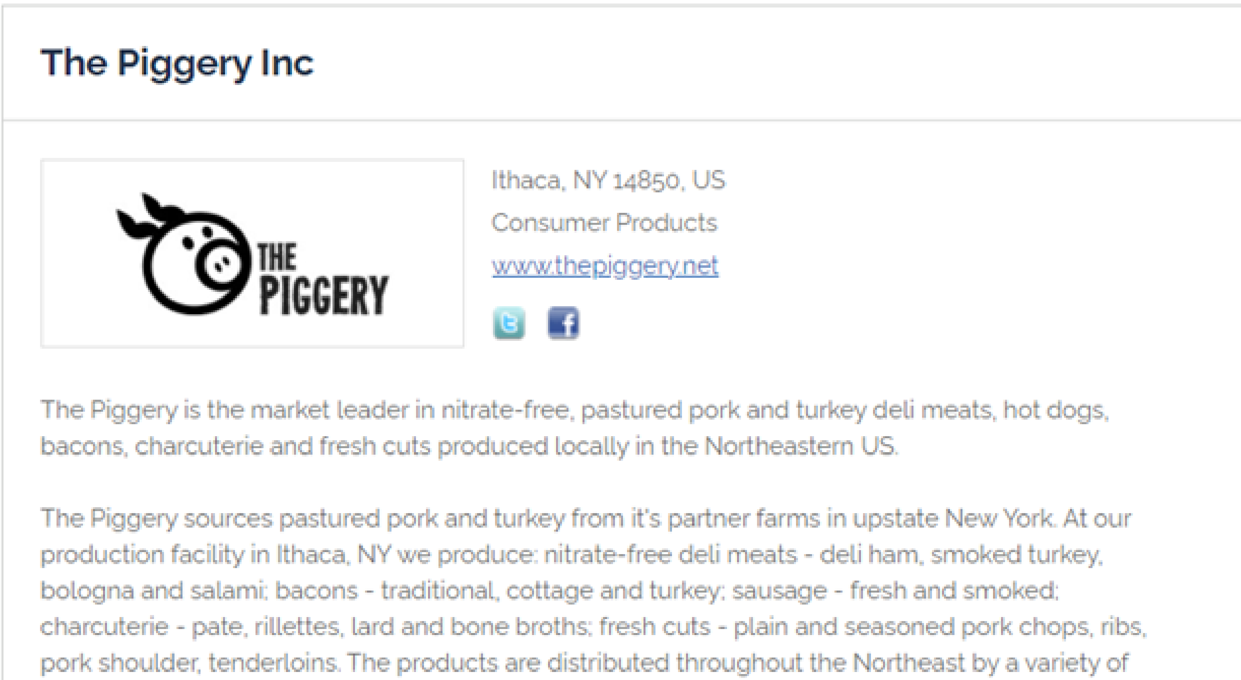

(a) Panel A

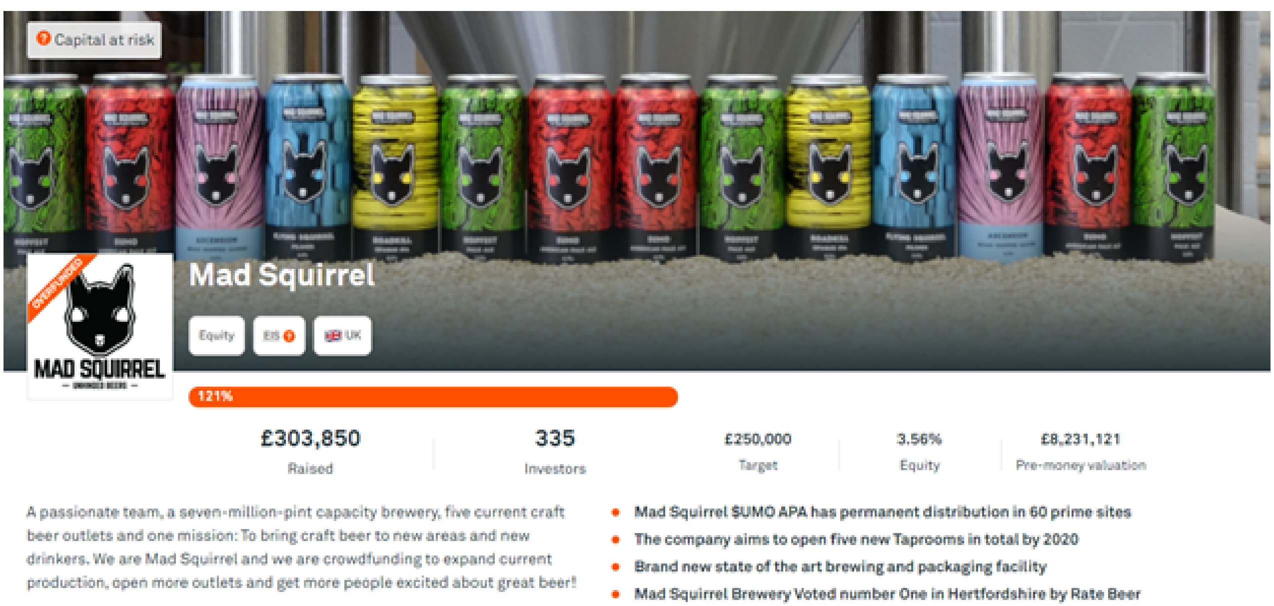

(b) Panel B

Source: Image from EquityNet.com The Piggery Mad-Squirrel

Figure 7: **Human Participant Results: Additivity** This reports the results of the human participation field study for informationally additive pitches on Crowdcube pitches.

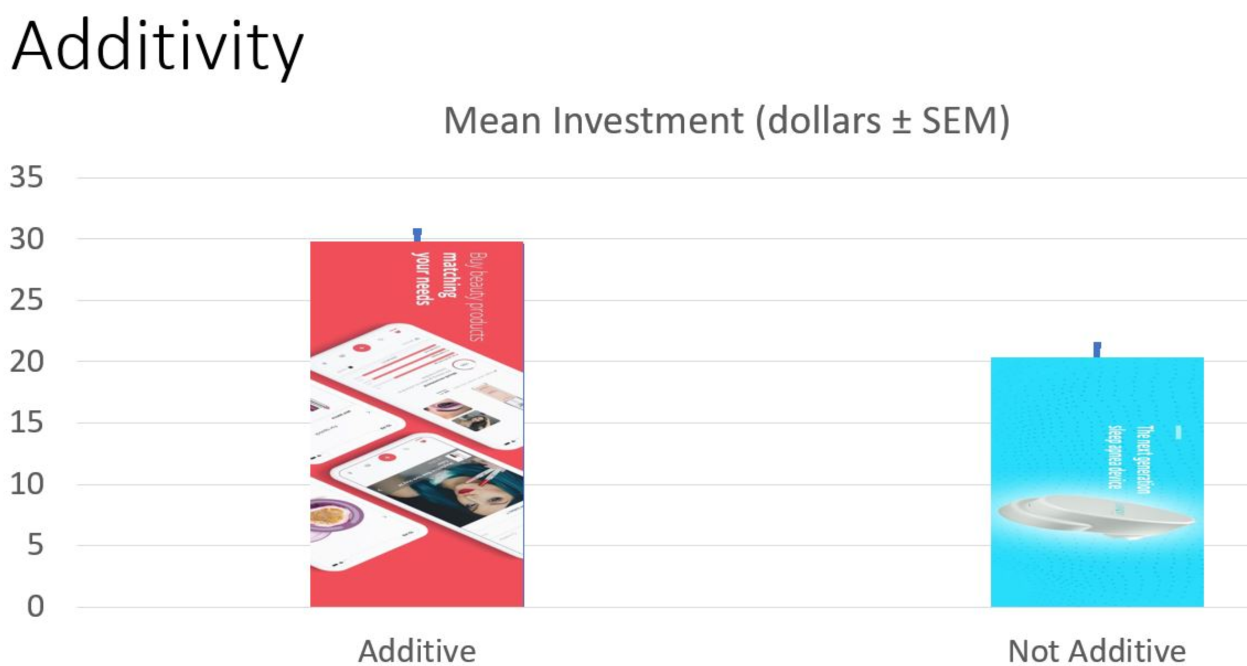

(a) Panel A

Source: Images from Crowdcube.com

Figure 8: **Human Participant Results: Intensity and Singularity** This reports the clarity results of the human participation field study based on Crowdcube pitches. Panel A presents the results for Singularity and Panel B presents the results for Intensity.

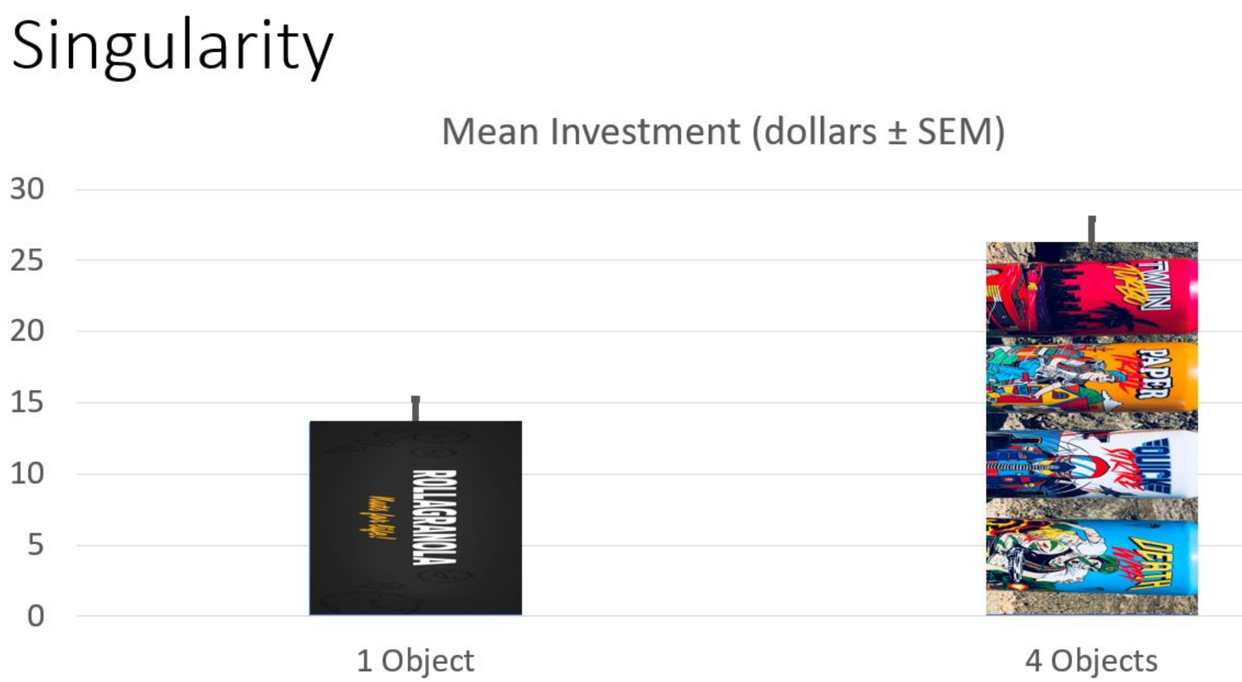

(a) Panel A

(b) Panel B

Source: Images from Crowdcube.com

Figure 9: **Human Participant Results: Object Discernibility and the Interaction Between Singularity and Intensity** This reports the clarity results of the human participation field study based on Crowdcube pitches. Panel A presents the results for Object Discernibility and Panel B presents the results for Interaction Between Singularity and Intensity.

(a) Panel A

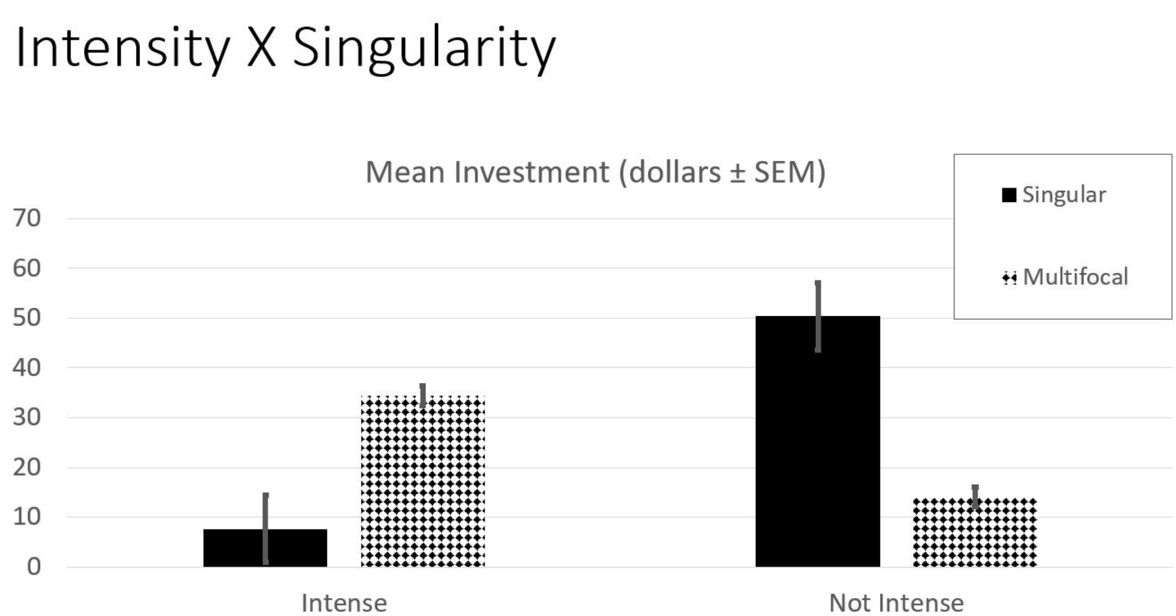

(b) Panel B

Source: Images from Crowdcube.com