Received: 14 March 2022 | Revised: 17 August 2024 | Accepted: 25 August 2024

DOI: 10.1002/smj.3665

RESEARCH ARTICLE

SMS | Strategic Management Journal WILEY

Are accelerators akin to breweries or wineries? A Bayesian variance decomposition of accelerator and cohort effects

Gil Avnimelech[1] | Gary Dushnitsky[2] | Florian Ellsaesser[3] | Markus Fitza[3,4]

[1]Ono Academic College, Kiryat Ono, Israel

[2]London Business School, London, UK

[3]Frankfurt School of Finance & Management, Frankfurt, Germany

[4]Business School, University of Exeter, Exeter, UK

Correspondence
Gary Dushnitsky, London Business School, London, UK.
Email: gdushnitsky@london.edu

Abstract

Research Summary: The literature on startup accelerators uncovers multiple factors associated with accelerators' advantages. Yet, we have a limited understanding of the relative magnitude of these factors. We ask: Are accelerators akin to breweries, where quality is mainly a function of the institution of origin (i.e., brewery for beer, accelerator for startups); or are they similar to wineries, where quality varies across cohorts (i.e., for a given winery, some vintages are of higher quality)? We explore this question using data from 1,350 tech-startups graduating from dozens of accelerators in a global technology hub. A Bayesian hierarchical variance decomposition approach is introduced to account for the highly-skewed zero-inflated distribution in startups' performance. We find that a notable fraction of startup performance is due to vintage; within-accelerator, cross-cohort variation.

Managerial Summary: Startup accelerators (i.e., short-term programs designed to help startups grow) are highly popular, with dozens of accelerators operating around the globe. Our focus is on accelerator programs aimed at catapulting technology ventures

The authors contributed equally and are listed alphabetically.

towards high growth. We ask: Are accelerators akin to breweries, where quality is mainly a function of the institution of origin (i.e., brewery for beer, accelerator for startups); or are they similar to wineries, where quality also varies across cohorts (i.e., for a given winery, some vintages are of higher quality)? A Bayesian hierarchical variance decomposition approach isused to study data from a global technology hub, detailing the performance of hundreds of startups that graduated across multiple accelerators. We find that a significant portion of startup success is linked to cohort-specific factors within accelerators, highlighting the role of timing and dynamics of each accelerator cohort.

KEYWORDS

accelerators, Bayesian hierarchical model, entrepreneurship, scale-ups, variance decomposition

1 | INTRODUCTION

Since the launch of Y Combinator's first cohort of eight startups in 2005, startup accelerators have become a well-known feature of the entrepreneurial-finance landscape. We define accelerators as entities that run fixed-term, cohort-based, business-support programs for nascent startups. They provide mentorship, educational, and networking input, often in return for an equity stake in the startup. About one-third of all startups in the United States that raised at least one round of VC funding in 2015 had been through an accelerator program, according to data provider Pitchbook (Mikey, 2016). The phenomenon is not limited to the United States: London witnessed its first major accelerator cohort in 2008 and now boasts dozens of accelerators (Dushnitsky & Sarkar, 2022; Miller & Bound, 2011). Accelerators are also popular in other countries, such as Chile (González-Uribe & Leatherbee, 2018) and India (Sharma et al., 2014). One of the largest pools of startups and accelerators outside the United States is in Israel, also known as Startup Nation (Senor & Singer, 2011).

As a result, startup accelerators have been a subject of interest from entrepreneurs, investors, policy makers and scholars. Our literature review (below) reveals keen insights from dozens of studies on the topic. One stream of work documents several accelerators and compares the performance of graduating startups (Assenova, 2020; González-Uribe & Leatherbee, 2018; Thompson, 2005; Yu, 2020). Another stream takes a complementary approach, using qualitative and observational data to flesh out the dynamics within a handful of accelerators (Clarysse et al., 2015; Cohen, Bingham, & Hallen, 2019). Thus, we have a breadth of knowledge about multiple factors through which accelerators affect startups' performance (Table 1). The insights guide policy makers and investors who design and launch accelerators across the world. As the number of studies and distinct factors increases, the challenge is knowing which are the critical factors to heed.

TABLE 1 Theoretical mechanisms explaining variation in the performance of accelerator graduates: An overview of plausible mechanisms.

Source of variation Mechanism	Accelerator Top-down	Manager Top-down	Cohort Bottom-up
Learning	• Stronger curricula • More comprehensive curricula • Bespoke curricula • Applied hands-on training	• Stronger curricula • More comprehensive curricula • Bespoke curricula • Proactive follow-up of startup progress	• Learning as a function of physical, temporal and structural proximity with other startups • Increased pace of learning and sharing across participants • Feedback through constant benchmarking to peer • Tournament effect • Spillover effects
Mentorship	• More mentors • More experienced mentors • Better mentor-founder matching • Mentors' training and guidance	• More mentors • More experienced mentors • Better mentor-founder matching • Mentors' training and guidance	• Startup-peer mentorship • Sharing knowledge of resources, talent and best practices
Networking	• More events • Better events	• More events • Better events	• Access to the existing network of all cohort participants • Thriving community during accelerator program • Life-long sense of community and networking, post-accelerator program
Certification	• Graduating a successful "unicorn" • Graduating multiple unicorns • Number of past cohorts • PR activities	• Within the manager's professional network • PR activities	• Being associated with other successful startups in a given cohort

The purpose of our study is to address this challenge. We offer a way to structure existing findings and guide future work. Our contribution to the accelerator literature is in (i) organizing the theoretically distinct factors into a cohesive framework and (ii) empirically documenting their relative impact. Building on prior work, we observe that there is a plethora of mechanisms through which a startup benefits from an accelerator experience. We organize these factors into a cohesive framework, distinguishing broadly between top-down mechanisms (i.e., mechanisms that are decided and designed by the accelerators and their managers) and bottom-up mechanisms (i.e., those emerging due to cohort dynamics).

We proceed to ask the following research question: To what degree are accelerators akin to breweries, where outcomes vary mainly as a function of origin (i.e., brewery for beer; accelerator for startups), and to what degree are they similar to wineries, where outcomes differ greatly across vintages (i.e., for a given winery, certain vintages may be of higher quality, and for a given accelerator, certain cohorts may exhibit superior performance)? By tackling this question, we can understand what is driving the performance of accelerator graduates. The goal is to uncover the relative impact of top-down mechanisms that are at the discretion of an accelerator and its managers (e.g., curricular choices) and bottom-up mechanisms where a cohort transforms into a vibrant community (e.g., sharing of information or vicarious learning from others). We recognize that there are excellent studies that examine these mechanisms. Our paper joins this conversation; rather than introducing new mechanisms, we examine the relative impact of each factor.

We undertake a variance-decomposition analysis to that end. The analysis is underpinned by an empirical and a methodological contribution. Empirically, we constructed a dataset of the universe of Israel-based accelerators and the startups that attended them. The dataset has three advantages. First, Israel is home to one of the most vibrant entrepreneurial ecosystems globally (Avnimelech & Teubal, 2006; Engel & del-Palacio, 2011; StartupBlink, 2021). Israeli startups are a major source of global innovation, from navigation (Mobileye, Waze) through e-commerce (ICQ, Wix) to semiconductors (Primesense, Anobit). Second, and relatedly, the setting complements accelerator studies, which are based on detailed data from the United States (Hallen et al., 2020; Yu, 2020) and Latin America (González-Uribe & Leatherbee, 2018). Third, the data are both comprehensive and fine-grained: we cover 1350 startups and have detailed data on the cohorts they were a member of as well as the identity of the accelerator managers who led them.

The choice of a variance-decomposition methodology is guided by seminal works in the field of strategy (McGahan & Porter, 1997; McGahan & Victer, 2010). This approach has been employed to move research forward by uncovering the relative importance of effect classes. Put simply, variance-decomposition studies serve as a "call to action" whereby the findings highlight the main sources of variation and thus direct scholarly attention and future work. The application of variance decomposition necessitated careful consideration (Vanneste, 2017). Accordingly, we introduce a Bayesian hierarchical model to model the dependent variable that is uniquely fitting for our setting (Gelman et al., 2013). We first detail the key departures from commonly studied settings; for example, extant strategy studies conduct variance decomposition for large publicly-listed firms where data are abundant and outcomes follow a normal distribution. In contrast, our setting concerns a smaller numbers of budding startups (which can result in sparse data), and outcomes follow a zero-inflated right-skewed distribution. The study proceeds to carefully explain the setup, estimation and presentation of the Bayesian hierarchical model and its results.

The analysis yields insights into the factors underlying accelerators' contributions and their relative impact. For example, extant work finds that the success of accelerator graduates is associated with cross-accelerator variation in observable factors. Our variance decomposition reveals that a substantial fraction of startups' performance is due to within-accelerator factors; namely, due to cohort vintage. Notably, the cohort effect (i.e., capturing performance differential among startups that are due to cohort dynamics) is of a larger magnitude than the accelerator effect (i.e., which captures startup performance differences across accelerators). This observation indicates the importance of understanding bottom-up cohort dynamics in addition to the well-studied topic of top-down mechanisms. This insight might be familiar to scholars from their own teaching experience: while every cohort of an accelerator (or a stream of a course) may be privy to the same top-down training and resources (syllabus and instructor), it is the case that one or two cohorts (individual classes) become a cohesive and thriving community, while the others do not.

The study makes several contributions. First, it derives theoretical implications to inform the accelerator literature. One implication of the fact that outcomes vary substantially within accelerators is that accelerators are, to some degree, like wineries: cohort vintage matters. A second implication is that the outcomes are not solely driven by top-down design choices. Rather, they may be due to bottom-up dynamics (e.g., the vintage effect arises at least partially because some cohorts transform into vibrant communities). These theoretical implications underscore the opportunity for further work on cohort dynamics (Assenova, 2020; Cohen, Bingham, & Hallen, 2019; Hallen et al., 2020). Third, we introduce a methodological approach—Bayesian hierarchical modeling—that is useful for studying variance decomposition of entrepreneurial outcomes, or other settings characterized by highly-skewed outcomes or sparseness across the hierarchical structure. Finally, the analysis is based on a dataset of the universe of accelerators for technology-based growth-orientated startups in Israel. It thus complements our knowledge of accelerators for growth-orientated startups that draws predominantly on the United States. The data are also unique in that it offers fine-grained information not only on the accelerators but also the managers who run them.

2 | DEFINITION

The accelerator literature is voluminous, and there is a similarly large number of definitions of "accelerator." Earlier work defined accelerators in comparison or in contrast to the wider business incubation literature. Along these lines, we see definitions of accelerators as "a new incubation model" (Clarysse et al., 2015), or "an emerging incubation-like model" (Yang et al., 2018). Notable points of differentiation from the traditional incubation model were the shorter duration within a program cohort—usually 3–6 months—and the fact that accelerators provide critical resources to the participating startups. More recent work proceeded to define accelerators not in contrast to an existing model, but rather based on their innate attributes (Cohen & Hochberg, 2014). We also witness work that focuses on differences across accelerators in terms of their backers (e.g., corporate vs. university accelerators), or industry focus (e.g., social accelerator, ecosystem builder) (Cohen, Bingham, & Hallen, 2019; Mansoori et al., 2019; Prexl et al., 2019).

We define accelerators as entities that run fixed-term, cohort-based, business-support programs for nascent startups. They provide mentorship and educational and networking input; often in return for an equity stake in the startup. A typical feature of accelerators is that they admit a cohort of startups per intake; for example, about a dozen startups per cohort (Hallen et al., 2020; Yu, 2020). The participants usually receive office space (Clarysse et al., 2015; Clarysse & Yusubova, 2014; Drori & Wright, 2018; González-Uribe & Leatherbee, 2018; Thompson, 2005), help with product development (Avnimelech & Rechter, 2023; Cohen, Bingham, & Hallen, 2019; Crişan et al., 2021), financial and legal support (Clarysse et al., 2015; Crişan et al., 2021; Glinik, 2019; Thompson, 2005; Uhm et al., 2018; Yang et al., 2018); HR/recruitment support (Banc & Messeghem, 2020; Lall et al., 2013), help with technical issues (Radojevich-Kelley & Hoffman, 2012), and networking opportunities (Avnimelech & Rechter, 2023; Cohen, Bingham, & Hallen, 2019; González-Uribe & Leatherbee, 2018; Kohler, 2016; Wright et al., 2017). Many accelerators focus on the personal development of the founders, equipping them with entrepreneurial skills, knowledge, self-efficacy, and legitimacy (Avnimelech & Rechter, 2023; Bischoff et al., 2020). Another feature of the accelerator experience is fundraising support. Accelerators host a "Demo Day" which is a public pitching event. It marks the culmination of the accelerator experience as a cohort of graduating startups pitch to investors (Clingingsmith & Shane, 2018; Cohen & Hochberg, 2014; Dushnitsky & Sarkar, 2022).

An important observation is that accelerators deliver these benefits with limited staff. The accelerator literature often makes note of the mentors who work with the startups or the venture partners who are usually involved with the screening process. Yet, these roles are usually undertaken by individuals for whom the accelerator is not their main professional responsibility, nor is it their main source of employment or income. In fact, the average accelerator has less than a dozen full-time employees (Cohen, Fehder, et al., 2019; Rechter & Avnimelech, 2024). Key personnel consist of the accelerator manager, along with additional employees who work in specific support roles (such as providing marketing and PR for the accelerator or engaging in various technology and business support to the startups). Notably, accelerator managers vary not only in their personal characteristics but also in the role they undertake. In most accelerators, managers play an important role in choosing which startups to admit into the accelerator. Rechter and Avnimelech (2024) further observe that managers undertake somewhat different roles. Some operate as project managers (e.g., coordinating the entire process); others play the role of a mentor for specific startups, or a super-mentor (i.e., the professional authority for all the mentors), or that of a process mentor.

3 | THE UNDERLYING MECHANISMS

Accelerators afford a window on entrepreneurial resource mobilization in their earliest phases (Clough et al., 2019). That said, the underlying individual and organizational mechanisms are often difficult to observe, especially over time and across multiple accelerators. Fortunately, there is a growing body of qualitative and quantitative work that sheds light on the dynamics within a focal cohort. These studies expose mechanisms that could be at play in other accelerators and cohorts.

3.1 | Between-accelerator differences: The role of top-down mechanisms

We first review key accelerator studies documenting the mechanisms decided and designed by the accelerators. We refer to these as "top-down" mechanisms because they are driven by the accelerators and/or their managers. Broadly speaking, the accelerator experience offers startup founders with three distinct opportunities (Drori & Wright, 2018).

First, it offers learning opportunities. The accelerator equips founders with the technical and business skills necessary to develop a successful startup. During their time in the accelerator, cohort members go through a pre-designed curriculum that covers strategic, technical, financial, and legal aspects of a business. They are also exposed to commercial and sales practices. These learnings are useful in searching and validating product-market fit, winning and scaling sales, and securing intellectual property. In this respect, rich qualitative data from US-based accelerators reveal that learning is a major mechanism impacting the success of graduating startups (Hallen et al., 2020). Specifically, accelerators help startups to identify the areas that are most critical to their future success and effectively engage in further information search (Cohen, Bingham, & Hallen, 2019). Quantitative work offers consistent evidence. González-Uribe and Leatherbee (2018) compare the provision of capital to the provision of schooling plus capital within the Startup Chile accelerator. They find that the latter is associated with performance 5 years after graduation. Startups' success is attributed to the skills and know-how acquired during their time at the accelerator (González-Uribe & Reyes, 2021). Learning how to make progress through experimentation, also known as “lean startup,” is one such skill (Avnimelech & Rechter, 2023; Camuffo et al., 2020; Leatherbee & Katila, 2020; Mansoori et al., 2019; Shankar & Clausen, 2020).

Second, the accelerator experience is designed to offer multiple opportunities to engage and network with external stakeholders. The most salient example is the accelerator mentors. Each startup is assigned a mentor or a group of mentors who work closely with the founders on startup-specific issues. The mentors not only expose founders to best practices, but also work with them in applying tools and frameworks specific to the startup. Moreover, mentoring is particularly valuable for entrepreneurs with limited pre-entry experience (Assenova, 2020). An additional benefit of being assigned a mentor is the opportunity to draw on the mentor's social networks. The accelerator experience exposes entrepreneurs to other stakeholders as well. For example, accelerators often cultivate a network of strategic partners at major corporations that enable participating startups to explore sales leads and develop business partnerships (Avnimelech & Rechter, 2023; Drori & Wright, 2018). Another example pertains to networking with investors during the pitching Demo Day (Dushnitsky & Sarkar, 2022). Finally, for long-standing accelerators, there is also networking with founders who graduated from past cohorts.

Third, accelerator affiliation bestows certification. Admission into a prestigious accelerator serves as an important signal in and of itself (Avnimelech & Rechter, 2023). Recall that early-stage startups usually have limited social proofs (e.g., sales traction, patent awards) (Stuart et al., 1999). At that early stage of development, admission into an accelerator serves as one of the few visible cues of quality. It can help the startup differentiate itself and attract prospective investors, customers, and employees. At the extreme, a startup may benefit from attending an accelerator—even if it gained no learning or social capital—simply through the certification effect.

The top-down decisions shaping the learning, networking, and certification opportunities are in the hands of the accelerator manager. We know that variation in inter-firm performance is in part due to changes in firm leadership (Fitza, 2014; Hambrick & Quigley, 2014; Quigley & Graffin, 2017). Similarly, the performance of accelerator graduates may be affected by changes to the individuals who manage the accelerator. Different managers bring different curriculum foci, and are likely to engage and mobilize different networks. Managers also decide which startups to admit into the accelerator programs; they might have different preferences and/or different abilities to spot startups that can benefit most from participation. Stated from a theoretical perspective, under different managers one might observe that different graduating startups are endowed with different types or magnitudes of skills and social capital. It follows that variations in startups' performance are, in part, due to the characteristics of the managers and the design choices they undertake. In this respect, the manager effect is similar to the aforementioned studies of accelerator effects, which attribute graduates' performance to top-down design choices. To return to our brewery metaphor, the accelerator manager can be seen as akin to the brew master, who might make adjustments to the beer recipe based on their specific knowledge or past experience.

The discussion highlights the role of accelerator and manager-led mechanisms. Past work documented the learning, networking and certification benefits. The scale of benefits varies by accelerator and manager. Stated in terms of variance decomposition, the degree to which the performance of startups graduating from accelerators differs depends on which accelerator they were a member of and who managed the accelerator at the time.

3.2 | Between-cohort differences: The role of bottom-up mechanisms

In addition to such top-down mechanisms, the impact of accelerators may also vary due to bottom-up mechanisms. It can be argued that startups' performance is shaped, to a certain degree, by bottom-up processes and the dynamics of a given cohort. This observation may resonate with scholars' personal teaching experience: we often observe that two streams of a course can exhibit very different dynamics. Returning to the winery analogy, the quality of a given bottle of wine is shaped not only by the vinery it originated from, but also by its vintage. Even a reputable winery may yield a less successful vintage for reasons that are beyond the control of its managers. Similarly, some of the benefits of being part of an accelerator could be specific to the focal cohort in which a startup partakes. If bottom-up processes are impactful, we would observe that cross-cohort variation—even within a focal accelerator and under the same manager—explains a notable fraction of variation in startup outcomes.

To the extent that bottom-up processes are impactful, a cohort becomes more than a collection of individuals; it forms a community. It carries critical implications to the direction of scholarly attention. If the cross-cohort variation is of meaningful magnitude in comparison to that observed across accelerators, it calls for scholarly investigation of bottom-up processes. This observation is informed by the social-psychology literature. The literature suggests that bottom-up processes involve the intense proximity and interactions, and activities that unfold throughout the program. Hence, we shift to discussing a set of within-accelerator, cohort-specific mechanisms. These include the intensity of collaboration and competition among the participants, and the vicarious learning that unfolds. Such a cohort effect calls attention to a different set of theoretical mechanisms; for example, learning by observing others' contemporaneous successes and failures (versus learning from one's own past success and failure), and peer-to-peer sharing (versus top-down instructor or mentor teaching), which can result in a thriving community of practice.

The social-psychology literature offers some guidance regarding the distinct theoretical mechanisms associated with such bottom-up cohort effects. Specifically, social-identification theory argues that proximity (physical, temporal, and structural) leads to information sharing and learning (Huang et al., 2013). These insights apply to accelerator cohort members who spend an intense period under conditions of high physical, temporal, and structural proximity. Each accelerator cohort comprises of several founders who often work within the same physical business premises: the accelerator space. This typically takes the form of an open-plan office, similar to a co-working space, where the entrepreneurs work shoulder-to-shoulder, occupying contiguous workstations. Recent evidence of learning benefits from physical proximity in similar entrepreneurial spaces (Roche et al., 2022) suggests that founders could experience enhanced learning from the direct physical interactions throughout the duration of the accelerator program.

The cohort experience is also characterized by temporal proximity. Each accelerator cohort has a clear start and end date and follows a pre-defined schedule to which entrepreneurs adhere. Simply put, the participants are usually in the same place at the same time. Finally, entrepreneurs often exhibit high structural proximity. That is, their startups are usually at a similar point in the life cycle and, therefore, face similar business challenges. Moreover, the accelerators often follow a structured curriculum in which all cohort members are given similar tasks at about the same time. Taken together, these features suggest that cohort peers face similar challenges at about the same point in time and therefore can, and will, draw on each other for information, advice, and referrals.

Recent accelerator studies document bottom-up cohort effects. A study of US-based accelerators finds that frequent and intense interactions with cohort peers is a key contribution to the accelerator experience (Hallen et al., 2020). Evidence from a Chinese setting reveals that monthly meetings with peer firms are associated with increased performance as firms engage in information sharing and mutual learning (Cai & Szeidl, 2018). Similarly, a randomized trial in the United States finds that constant interaction around common social interests facilitates information flows (Krishnan et al., 2021). Another randomized trial finds that cohort members themselves are an invaluable source of advice regarding people management (Chatterji et al., 2019).

Not every cohort becomes a thriving community, as proximity can also have disadvantages. At times, collaboration may be overtaken by competition. This is especially the case among entrepreneurs who are spatially, structurally, and temporally proximate. Extant work documents evidence of adverse competitive dynamics among cohort members. The negative effects arise due to two distinct mechanisms: (i) spillover concerns and (ii) benchmarking and envy. Proximity can give rise to advantageous learning but can equally be the source of valid spillover concerns. The concerns pertain to the leakage of technical, commercial or strategic information. Leakage may be due to deliberate action as one cohort member expropriates from another, but could also take place unintentionally among the co-located startups. It is not surprising to observe adverse competitive dynamics among those targeting related product markets (Cai & Szeidl, 2018). Benchmarking is another mechanism that can impede collaboration. There is evidence of intense within-cohort competitiveness when the interactions are of a tournament nature (Krishnan et al., 2021). This is consistent with the concept of social comparison that underlies envy and which has a well-known effect on organizational dynamics and outcomes (Nickerson & Zenger, 2008). The spatial, structural and temporal proximity in which cohort members operate makes comparisons to other members highly salient. This effect may arise irrespective of whether the accelerator adopts formal benchmarking practices. It can give rise to envy and exacerbate bottom-up competitive dynamics.

The discussion suggests that some cohorts turn into a thriving community. Indeed, a cohort features the characteristics of many communities of practice where members share interests and interact with boundaries of space and time (Wenger, 1998). Qualitative data find that entrepreneurs benefit from the energy, friendship and support co-working entails (Howell, 2022). That said, the community-of-practice also exhibits variation in the sense that not every cohort blossoms into a community. Even when members share similar structures, bottom-up dynamics can lead to substantial variation in the flow of knowledge among group members and, thus, the value they derive (Probst & Borzillo, 2008; Thompson, 2005; Wenger et al., 2002).

To conclude, we highlight pathways through which the accelerator experience impacts startups' performance. These include accelerator-led and manager-led top-down mechanisms as well as bottom-up cohort dynamics. Table 1 summarizes the mechanisms underlying these effects. Some of them have already been the subject of scholarly work (including work documenting several bottom-up dynamics; Cohen, Bingham, & Hallen, 2019), while others could be explored by future work.

What is to be gained by differentiating between accelerator, manager and cohort effects on startup performance? Identifying the source of performance differences is important because the differences can inform different predictions. Consider, for example, the effect of cohort industrial homogeneity (where all participating startups target the same industry). From an accelerator perspective (rightmost column in Table 1), there are a number of top-down mechanisms that should result in better startup performance. By focusing on a particular industry, the accelerator can develop more robust and targeted curricula, and thereby facilitate better learning. In contrast, a cohort perspective (leftmost column in Table 1) may result in less conclusive predictions. The structural proximity associated with same-industry participants may be designed, top-down, to encourage learning. At the same time, we recognize that the sharing of experiences and advice among peers constitutes a bottom-up dynamic that may be curtailed where participants view each other as direct competitors. In other words, an industry-focused accelerator may be designed to increase the potential for top-down learning benefits, but that potential may fail to materialize due to adverse bottom-up dynamics within the cohort.

While there is work that examines each of these cohort mechanisms, our study asks a different question. It looks at the pathways in a holistic manner. As a variance-decomposition study, our focus is not on identifying the underlying mechanisms behind each effect, but on examining the relative size of each effect. To paraphrase Withers and Fitza (2017), such an analysis answers a fundamental question: If we want to understand how startups benefit from accelerators, how much does studying the accelerator itself, its manager, or the cohort, help toward this goal? We believe that addressing this question carries critical implications to the direction of scholarly attention. If the cross-cohort variation is of similar magnitude to that observed across accelerators, it calls for scholarly investigation to equally recognize (and study) bottom-up processes as it does top-down accelerator structures and management.

4 | DATA

Our study uses a hand-collected dataset of the universe of Israel-based accelerators and the startups that attended them. The setting complements detailed studies of US-based accelerators (Hallen et al., 2020; Yu, 2020). Israel is home to one of the most vibrant entrepreneurial ecosystems, and many entrepreneurs and investors there have close links with their counterparts in the United States and Europe (Avnimelech & Teubal, 2006; Engel & del-Palacio, 2011; StartupBlink, 2021).

The data-collection effort proceeded as follows. First, we were mindful that many organizations present themselves as "accelerators." Per prior work, we focused on those entities that met the following criteria: (i) they operate a cohort-based program; (ii) the program is fast-paced (i.e., 9 months' duration or shorter); and (iii) the program includes an educational component and mentorship. We excluded a few short-lived accelerators that had seven startups or fewer. Second, we identified every startup that graduated from an accelerator during the period 2010–2019. It is a period that has seen substantial startup and accelerator activity. It is also a meaningful period in the sense that it sits between two major events; the Great Financial Crisis and the Covid pandemic. To construct the data, we triangulated three different sources, including a leading startup directory (IVC Research Center Israel; Conti, 2018), the websites of the accelerators, and LinkedIn profiles of entrepreneurs and accelerators' managers. We further used this information to confirm the identity and cohort of each startup and the identity of the accelerator manager of each cohort.

These steps yielded a dataset of all accelerator graduates in Israel during the 2010s. It includes the following variables: startup name, name of the accelerator it graduated from, cohort start and end dates, name of the accelerator manager at the time, and startup industry.[1]

[1]The IVC defines the following industries: Cleantech, Communications, Fintech, Life Science, Semiconductors, Social/Impact Ventures, Internet, Software, Other Technologies.

TABLE 2 Dimensions of sample.

Sample	Startups 12 months after entry				Startups 3 years after entry			
	Total obs.	Avg. per accel.	Avg. per manager	Avg. per cohort	Total obs.	Avg. per accel.	Avg. per manager	Avg. per cohort
No. of accelerators	24				17			
No. of managers	69	2.88			43	2.53		
No. of cohorts	158	6.58	2.29		89	5.24	2.07	
No. of startups	1350	56.25	19.57	8.54	515	30.29	11.98	5.79

We further collected information about the startups' features (e.g., date of incorporation, the total amount of capital raised prior to entering the accelerator) and track record (e.g., startup's current status, the amount raised after joining the accelerator). Data were collected through October 2020.[2]

Our study aims to unpack the accelerator, manager and cohort. Accordingly, we exclude cases where the effects are indiscernible. To avoid confounding the cohort effect with individual startups, we removed "cohorts of one" (i.e., an accelerator cohort that only had one member). We similarly excluded accelerators with less than two managers and all managers with less than two cohorts, because in these cases the manager effect cannot be separated from the accelerator effect, and the cohort effect cannot be separated from the manager effect.

These efforts resulted in detailed data for 1350 startups graduating from 158 cohorts across 24 accelerators that were managed by 69 different managers. We believe that the database constitutes one of the notable contributions of our study because it complements other large-sample analyses that are often based on the GALI data. The two databases complement each other as they cover a similar period (GALI covers 2013–2019), with our sample focusing on technology-based growth-oriented startups, and GALI including many startups and accelerators geared toward socially-orientated employment-enabling goals.[3] Table 2 summarizes the key dimensions of the sample for each of the dependent variables (DVs) (see below).

5 | METHODS

We conduct a variance-decomposition analysis to assess the magnitude of different effect classes (i.e., accelerator, manager, cohort, year, and industry) on startup success. Variance-decomposition analysis holds a notable place in the strategy field (Vanneste, 2017), with several studies credited for launching important streams of literature (e.g., McGahan & Porter, 1997; Misangyi et al., 2006; Quigley & Graffin, 2017; Schmalensee, 1985). Notably, past work utilized several different methodologies to conduct a variance-decomposition analysis, with the two dominant

[2]To ensure data availability we excluded startups that entered an accelerator after March 2019.

[3]The GALI website explains "*Why Accelerators? Since 2005, hundreds of accelerators have been launched … Investors, development agencies, and governments are excited by their potential to drive growth, spur innovation, solve social problems, and increase employment opportunities in emerging markets*" (www.galidata.org/about/).

methods being ANOVA (e.g., Fitza et al., 2009; McGahan & Porter, 1997; McGahan & Victer, 2010; Schmalensee, 1985) and linear multi-level modeling with maximum likelihood estimation (MLM) (see Vanneste, 2017 for a review). Each methodology brings its own strengths and shortcomings. It follows that a scholar should explicitly consider the nature of the underlying data and select an appropriate methodology. This is important due to the unique features of our analysis: we study fledgling entrepreneurial ventures, whereas extant work has predominantly focused on established corporations. Thus, we start with a detailed discussion of the nature of our data.

5.1 | Nature of the data

Extant variance-decomposition studies in the strategy literature focus on established corporations. As a result, most studies deal with companies for which abundant financial and accounting data are accessible through public databases such as Compustat. Given the established nature of the companies and the availability of the data, extant work usually studies firms' profitability using measures such as return on assets (ROA). Moreover, these studies benefit from large sample sizes as regulatory filing requirements imply that Compustat covers thousands of companies over a period of decades and across numerous industries.

We study startup accelerators that are an important source of innovations and value in the modern economy. As we undertake this path, we are careful to explain how our setting differs from that of the typical variance-decomposition study, and the methodology we have adopted to address this. We start by explaining the key points of difference: (i) the challenge of using a proxy of early-stage performance; (ii) the zero-inflated right-skewed distribution of our chosen performance proxy: fundraising; (iii) the size of the sample; and (iv) its implication for potential data sparseness. We thus discuss a methodology that meets the features of our setting; a Bayesian hierarchical model (Gelman et al., 2013; Mullahy, 1986; Wibbens, 2019).

First, we underscore a key difference between measuring the performance of a fledgling startup and that of an established company. Extant work utilizes detailed financial and accounting data to construct performance measures such as ROA. While this measure may be available and relevant to established companies, it is not applicable to early-stage startups. These startups engage in costly development efforts and are, therefore, predominantly unprofitable. Accordingly, we turn to the accelerator and entrepreneurship literature in pursuit of appropriate performance measures. Extant work uses fundraising amount as a key indicator of initial entrepreneurial success in general (Guler, 2007; Hallen, 2008; Lerner, 1994), and successful accelerator participation, in particular (Dushnitsky & Sarkar, 2022; Hallen et al., 2020; Yu, 2020). Following these studies, we utilize fundraising success as a proxy of startup performance. Specifically, we define two dependent variables. The first dependent variable captures the fundraising amount (in millions of US dollars) during the 12 months following entry into an accelerator. This variable proxies for the impact of the accelerator on startups' immediate goal: raising early-stage capital to fund development and growth. We observe fundraising during that first year for 1350 unique startups. The second dependent variable explores startups' fundraising over a longer period, that is, the total amount raised in the 3 years after entering an accelerator. Because the latter focuses on startups that entered at least 3 years prior to our sample end date, the sample has fewer observations: 515 startups.

Second, we examine a related issue concerning the distribution of the dependent variable. In the case of established companies, the performance measures usually exhibit a normal distribution (Vanneste, 2017). In contrast, the nature of the entrepreneurial setting is such that only a few startups fully succeed in their development efforts; some startups secure no funding after entering an accelerator, others raise modest or moderate investment, and a few secure significant investment amounts (Crawford et al., 2015; Scherer et al., 2000). As a result, the funding amount does not follow a normal distribution. Rather, it exhibits a highly skewed distribution with a right skew, as captured in Figure 1. This carries implications for our choice of methodology because most estimation approaches, such as the prevalent multi-level modeling approach (Misangyi et al., 2006), are based on certain assumptions about the distribution of the dependent variable. This guided our decision to use a more flexible Bayesian hierarchical model, as we explain below.

Third, past variance-decomposition studies in the strategy literature usually utilized Compustat data and, as a result, have a very large number of observations. We observe hundreds of unique startups in one of the most globally vibrant entrepreneurial hubs. Yet, our sample size might be seen as modest relative to previous variance-decomposition studies, but that is not necessarily the case. A recent meta-analysis of firm and industry variance-decomposition studies identifies 41 studies in the strategy field, of which about 25% employ fewer than 1000 observations (Vanneste, 2017). This insight indicates that our sample size, while smaller than the typical strategy study, is not uncommon. That said, the size of our sample can raise the issue of data sparseness, which we discuss in the next point.

Fourth, variance-decomposition analyses study the magnitude of effects classes that are often hierarchically nested. For example, strategy scholars studied the "business-unit effect," which is nested within the overall "corporate effect." In other words, there is a hierarchical structure where the corporate effect represents higher-level cells, and the business-unit effect is the lower-level cells. This structure is a common feature of variance-decomposition data, yet it can give rise to sparseness under some conditions. To see that, let us turn to discuss our setting. We study five effects classes—year, industry, accelerator, manager, and cohort. As with other studies, these effects represent a nested structure. Each startup is a member of industry and of a cohort, while cohorts are nested in accelerator managers (i.e., the manager of that cohort), and

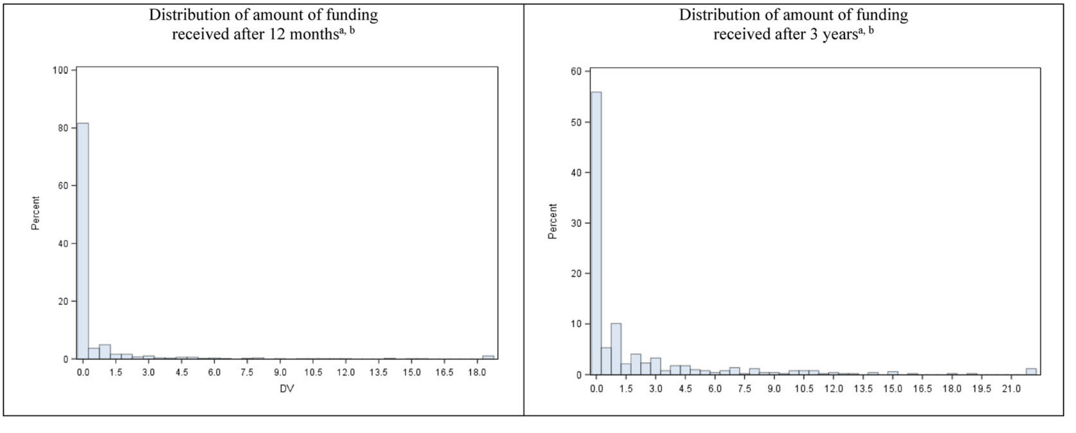

FIGURE 1 Distribution of the dependent variables. [a]For illustration purposes, large values are capped at the 99 percentile. [b]The scale on the x axis is $0.5 million. The high peek to the left represents all observations with a DV of 0 to 0.5. The percentage for observations with a DV of precisely 0 is 73% and 39% for the 12-month and 3-year DVs, respectively.

accelerator managers are nested in accelerators (i.e., the accelerator they are affiliated with). The accelerators and cohorts represent the highest-level and lowest-level cells, respectively.

In hierarchical models, the term "group size" refers to how many lower-level cells are contained in a higher-level cell. The issue of data sparseness arises when the group size becomes small. Prior work found that sparseness can result in inflated effect sizes (Stavropoulos et al., 2015; Vanneste, 2017).[4] Our sample size implies that we may face this issue. To that end, Table 2 summarizes the dimensions of our sample. It shows, for example, that in the 3-year fundraising variable, there is an average of 5.8 startups per cohort, 2.1 cohorts per manager, and 2.5 managers per accelerator. These values fall close to or below the "sparseness threshold" of five observations that Stavropoulos et al. (2015) determined for ANOVA-based variance-decomposition analyses.

Taken together, the four points above guided our methodological choices for executing the variance-decomposition analysis. The last two points suggest potential data sparseness. Detailed simulation analyses by Vanneste (2017) reveal that effect inflation due to data sparseness is a concern when using traditional variance decomposition methodologies, especially in the context of ANOVA analysis and, to a lesser extent, also in the context of multi-level modeling (MLM). We note that MLM and Vanneste's (2017) simulations, rely on normally distributed dependent variables, whereas our dependent variable follow another distribution (see points 1 and 2 above).

We, therefore, apply a different analytical methodology, a Bayesian hierarchical approach that meets the specific features of our setting. As we show below, the advantage of this approach is its high flexibility in defining the structure of the model in such a way that it fits our data characteristics (Gelman et al., 2013; Paez et al., 2008; Zhang & Liu, 2019). The next section explains the features of our main model and elaborates on how this methodology copes with the specific features of our sample.

5.2 | The Bayesian model

Our model is motivated by the characteristics of the data. First, as described above, our setting has a hierarchical structure with potential sparseness (e.g., with small lower-level group sizes). To see how we address this hierarchical structure with potential sparseness, it may help to understand the underlying challenge associated with estimating variability across sparse lower-level groups such as those represented by accelerator cohorts. Traditionally, a scholar can either (i) fit the data to an independent model for each of the groups (no pooling) or (ii) fit one model to the whole dataset across different groups (full pooling). In the context of sparse data, the first option, no pooling, of fitting independent models can suffer from having too few data points, which inhibits efficient estimates for every group and ignores information across groups (Gelman & Pardoe, 2006). The shortcoming of fitting a joint model, full pooling, is that it leads to a single parameter estimate for all groups, thus ignoring variability across groups.

The hierarchical structure of our Bayesian model effectively combines both approaches by using partial pooling, learning a prior distribution that collects the information that is valid across the groups and allows the parameters of each group to vary conditioned on this prior

[4]Stavropoulos et al. (2015) note: "In general, simulation studies that have examined the effect of sample and group size on variance estimates using mixed hierarchical and cross-classified models have shown that variance-decomposition analysis functions poorly [...] when group sizes are very small."

(Beaumont & Rannala, 2004). If, for example, only one observation is available for a group, the prior heavily constrains the estimate; that is, estimates for groups with small group sizes effectively use the data across all groups (Gelman & Pardoe, 2006; Korteweg & Sorensen, 2017). In contrast, the more observations are available in a group, the less is it the case that the estimation draws on the prior which is learned across all groups (Gelman & Pardoe, 2006).

Furthermore, the flexibility of the Bayesian model allows us to vary how prior distributions apply within each hierarchical level. For example, to learn the parameters for the joint distribution that the individual manager coefficients are drawn from, our model constrains the effect of managers on startups' funding by sharing the data of all the managers who managed a given *accelerator*, but we allow for the manager effect to vary *across accelerators*. In other words, the manager effect is not universal but can be different within different accelerators. The same is true for the cohort effect; it can differ for different managers.

The second important feature of our setting is the distribution of the dependent variables. Whereas the previous point concerns the structure of the data and the number of observations across the different effects classes (i.e., how many startups per cohort, how many cohorts per manager, how many managers per accelerator), the current point concerns the distribution of startup performance (i.e., fundraising post entry to an accelerator). The dependent variables follow a right-skewed distribution with many zeros (zero-inflated). To address this, our model uses a Tweedie distribution to model the dependent variables. Tweedie distributions are a family of probability distributions that is well suited to model data with many zeros and a long right tail (Dunn & Smyth, 2001; El-Shaarawi et al., 2011; Gilchrist & Drinkwater, 2000; Jorgensen, 1987; Jorgensen, 1997; Kokonendji et al., 2021; Tweedie, 1984; Ye et al., 2021).

Model specification: We estimate a model of a startup's fundraising success *after* it enters an accelerator. The general structure of our model follows previous variance decomposition works in similar hierarchical settings. We use a hierarchical model where startups are nested in industries and cohorts, cohorts are nested in managers which are themselves nested in accelerators. We use random intercepts for the hierarchical elements (industry, accelerator, manager, and cohort) and fixed-year effects (Misangyi et al., 2006; Quigley & Graffin, 2017; Withers & Fitza, 2017). We also include the funding a startup received before entering an accelerator (i.e., the variable *previous finding*) to control for inherent startup abilities that may affect our dependent variable (Conti et al., 2022; Guzman & Stern, 2016).

Our model is represented by Equations (1) and (2), where each startup is indexed by s, each accelerator is indexed by a, each accelerator has a set of managers indexed by m, and each manager manages a set of cohorts indexed by c.[5] In the model, the funds a startup raised after entering an accelerator are also affected by the industry of the startup (indexed i) and the year (indexed y) and by our aforementioned control variable (*previous funding*).

The model is estimated using a Bayesian hierarchical approach, with the dependent variable following a Tweedie distribution. Formally, y_s, the funding raised by startup s, after it joined accelerator a under manager m and cohort c is described by Equation (1).

$$y_s \sim Tw(\mu_s, \varphi, p). \tag{1}$$

This Tweedie distribution accounts for the zero-inflated and right-skewed distribution of the dependent variable (El-Shaarawi et al., 2011; Gilchrist & Drinkwater, 2000; Jorgensen, 1987; Jorgensen, 1997; Kokonendji et al., 2021; Tweedie, 1984). The parameters of the Tweedie distribution are the mean μ_s, the dispersion φ and the power parameter p, where each of the

[5]The bracket "(s)" in an index implies that the index is a function of the particular startup.

parameters takes non-negative values. The parameters φ and p together define the zero inflation and degree of right-skewedness of the distribution; these parameters are estimated by the model. Per prior work, we constrain p to be in a range between 1 and 2 which represents zero inflated distributions (El-Shaarawi et al., 2011; Kokonendji et al., 2021; Tweedie, 1984). The term μ_s represents the mean potential fundraising success of a startup after it entered an accelerator. It is defined by our regression estimation:

$$\mu_s = \beta_0 + \beta_{y(s)} * year_{y(s)} + \gamma_{a(s)} + \delta_{a(s),m(s)} + \varphi_{a(s),m(s),c(s)} + \theta_{i(s)} + \beta_s * previous\,funding_s + \varepsilon_s. \quad (2)$$

Here, the intercept β_0 represents the grand mean for the DV. The terms $\gamma_{a(s)}, \delta_{a(s),m(s)}, \varphi_{a(s),m(s),c(s)}$ and $\theta_{i(s)}$ represent random intercepts for accelerator, manager cohort and industry[6] (hierarchically nested as described above). $\beta_{y(s)}$ represents a fixed effect for the year startup s joined the accelerator ($year_{y(s)}$). β_s is the coefficient for *previous funding*$_s$. The error term ε_s captures startup-specific idiosyncratic performance differences.

The value of $\gamma_{a(s)}$ is the same for all startups graduating from a focal accelerator a. It thus captures the differences between startups that graduated from different accelerators. Depending on the contribution of its specific accelerator, a startup might perform better or worse.[7] Formally stated, our Bayesian model assumes that each accelerator receives an independent draw of $\gamma_{a(s)}$, which is drawn from a *prior* distribution $\gamma_{a(s)} \sim \mathcal{N}\left(0, \sigma_\gamma^2\right)$, and remains constant throughout an accelerator's life. The priors represent our beliefs about the "ability" of each accelerator to contribute to a startups' funding success. In line with Gelman et al. (2013, p. 113), we use a normal distribution for the priors as we have no knowledge about how accelerator abilities are distributed. Hence, our starting (prior) assumption is that the distribution of accelerators' contributions vary around a mean, and that contributions that are further away from the mean are less likely. This approach is consistent with Gelman et al. (2013); in the absence of knowledge about a given parameter, one can proceed by (i) assuming that the parameter follows a normal prior distribution and (ii) follow-up with post-estimation fitness checks to assess the validity of the assumption. To demonstrate the validity of this assumption, we can foreshadow that our post-estimation tests suggest the model offers a good fit.

The Bayesian model estimates a *posterior* of $\gamma_{a(s)}$ with a variance (σ_γ^2). Startups from accelerators with higher $\gamma_{a(s)}$ raised greater amounts, and startups with lower $\gamma_{a(s)}$ raised lower amounts after joining an accelerator. Thus, the variance of the posterior (σ_γ^2) reflects the degree to which startup fundraising success differs across accelerators. If the value of σ_γ^2 is low, there would be little variation in $\gamma_{a(s)}$ which implies that accelerators are similar in terms of how they contribute to the fundraising success of their graduates. In contrast, when σ_γ^2 the value of σ_γ^2 is high, there is substantial variation across accelerator graduates, which would suggest a notable "Accelerator Effect." That is, the fundraising success of startups after entering an accelerator is greatly affected by which specific accelerator they joined.

The variance σ_γ^2 (as well as the other posterior variances mentioned below) is the focus of our variance decomposition study. It represents the *accelerator effect*; namely, the degree to

[6]For *random intercepts*, we followed the convention to not write down separate variables and coefficients for simplification. Strictly speaking the effect for accelerator is $\beta_{a(s)} * \gamma_{a(s)}$, manager is $\beta_{a(s),m(s)} * \delta_{a(s),m(s)}$, cohort is $\beta_{a(s),m(s),c(s)} * \varphi_{a(s),m(s),c(s)}$, and industry is $\beta_{i(s)} * \theta_{i(s)}$.

[7]The accelerator effect represents how much accelerators vary around the average contribution associated with having been at an accelerator. Our focus is not with the average accelerator contribution (which is incorporated into the regression intercept) but rather the variation across accelerators, σ_γ^2.

which it "matters" which accelerator a startup was a member of. In other words, the term captures the extent to which affiliation with a different accelerator is associated with variation in startups' performance (i.e., post-accelerator fundraising success). The reported *accelerator effect* is a percentage calculated by dividing this variance by the model's total variance.

Similarly, $\delta_{a(s),m(s)}$ captures the *manager effect*. The regression coefficient is estimated based on a prior representing the impact of an individual manager above or below the average manager effect. The prior is drawn from a normal distribution $\delta_{a(s),m(s)} \sim \mathcal{N}\left(0,\sigma^2_{\delta,a}\right)$. We allow the variance of the distribution to vary across different accelerators, which is captured by the subscript a for $\sigma^2_{\delta,a}$. That is, managers within one accelerator might vary more (or less) than managers within another accelerator. The reported *manager effect* is calculated based on the average variance of the posterior distribution across all accelerators ($\sigma^2_{\delta,a}$).

Finally, we consider the *cohort effect*. The term $\varphi_{a(s),m(s),c(s)}$ captures the cohort effect with priors drawn from a normal distribution $\varphi_{a(s),m(s),c(s)} \sim \mathcal{N}\left(0,\sigma^2_{\varphi,m}\right)$. Similar to the manager effect, we allow the variance of that distribution to vary across managers, which is captured by the sub-script m for $\sigma^2_{\varphi,m}$.[8] As above, the reported *cohort effect* is based on the average variance of the posterior $\left(\sigma^2_{\varphi,m}\right)$. It reflects the degree to which startup fundraising differs across cohorts. Following that logic, the term $\theta_{i(s)}$ captures the industry effect with priors drawn from a normal distribution $\theta_{i(s)} \sim \mathcal{N}\left(0,\sigma^2_i\right)$.[9] More details about our model can be found in Appendix A.

The Bayesian hierarchical model is set up with weakly informative or noninformative hyper-priors (see Appendix A), so that the results are driven by the data rather than the priors (Gelman et al., 2013; Wibbens, 2019).[10]

Model estimation and post-estimation validation: We estimate our Bayesian model using a Hamiltonian Monte Carlo Sampler with 2000 iterations for warm-up followed by 2000 iterations and two Markov chains to simulate the posteriors (Gelman et al., 2013; Wibbens, 2019).

We also report a set of post-estimation validation tests. As described in Gelman (2013, p. 129) the definition of a Bayesian model is provisional until the model is validated in terms of its convergence and posterior distribution after it has been estimated. Accordingly, we extensively validated the model. We report the convergence of the Markov chains, effective sample size n_{eff}, $\widehat{R}$, warm-up and main posterior sampling (Gelman et al., 2013; Wibbens, 2019). Following Vehtari et al. (2017, 2015), we analyze the posterior distribution in terms of its mean, extreme value statistics (i.e., whether it represents large values and the zero-inflation well), and also present a dense overlay of posterior- and actual distribution and assess the distribution of the error term.

Finally, we conducted two tests to check the model's predictive performance in line with Gelman et al. (2013, chap 6) and Vehtari et al. (2015). The first test is the Pareto K-diagnostic leave-one-out cross-validation (LOO-CV) based on Pareto smoothed importance sampling (PSIS). And the second test is a posterior predictive checks (PPC) leave-one-out (LOO) cross-validation (CV) probability integral transformation (PIT) Plot combined with an outlier removed model based on Mangiola et al. (2021). Overall, all these post-estimation validations suggest that our Model addresses the special characteristics of the data well (see Appendix B for details).

[8]Because managers are nested within a specific accelerator, the subscript does not explicitly note the accelerator.

[9]The variance explained by the fixed year effect and by the control variable is calculated as per prior work (Misangyi et al., 2006; Matusik & Fitza, 2012; Quigley & Graffin, 2017; Withers & Fitza, 2017). Specifically, we calculate the difference in the explained variance for the model with and without the associated variables.

[10]In hierarchical models just as in non-hierarchical models, it is often practical to start with simple, relatively noninformative, prior distributions (Gelman et al., 2013, p. 108). This "reflects our ignorance about the unknown hyperparameters." The word "noninformative" indicates "our attitude toward this part of the model and is not intended to imply that this particular distribution has any special properties" (Gelman et al., 2013, p. 110).

TABLE 3 Mean percentage of variance explained.

Panel A: The dependent variable is capital raised 12 months after accelerator entry		
	Base model (accelerator effect only)	**Full model (accelerator, manager and cohort effect)**
Accelerator	8.92 (7.32–10.14)	3.60 (2.43–4.71)
Manager		7.70 (7.46–7.92)
Cohort		7.51 (7.37–7.68)
Industry	2.10 (1.01–2.86)	3.18 (1.06–4.55)
Year	0.00 (0.00–0.63)	0.00 (0.00–0.60)
Pre accelerator funding	0.09 (0.00–0.44)	0.07 (0.00–0.29)
Panel B: The dependent variable is capital raised 3 years after accelerator entry		
	Base model (accelerator effect only)	**Full model (accelerator, manager and cohort effect)**
Accelerator	8.93 (7.39–9.28)	2.60 (1.05–3.32)
Manager		4.91 (4.54–5.41)
Cohort		5.64 (5.24–5.99)
Industry	1.52 (0.81–2.57)	1.56 (0.43–2.31)
Year	0.00 (0.00–0.19)	0.00 (0.00–0.11)
Pre accelerator funding	0.00 (0.00–0.23)	0.00 (0.00–0.25)

Note: 95% posterior intervals in parentheses.

6 | RESULTS

Table 3 presents the main results. The table reports the mean percentages of the total variance in startup fundraising after entering an accelerator which is explained by each effect.

We begin by estimating a base model that contains the accelerator effect, as well as the year effect, the industry effect, and the pre-accelerator funding control. This facilitates comparison with past work and thus serves as a benchmark illustrating the importance of addressing the "brewery vs. winery" question. Panel A reports the results for the first dependent variable; the fundraising amount within 12 months of entering the accelerator. In this model, the mean accelerator effect is 8.9%, and the 95% posterior interval (PI), which is the Bayesian equivalent to the confidence interval, has a range of 7.3%–10.1%. The year effect is 0.0% (95% PI: 0.0%–0.6%), the industry effect is 2.1% (95% PI: 1.0%–2.9%), and the pre-accelerator funding control variable explains 0.1% (95% PI: 0.0%–0.4%) of the variance in post-accelerator fundraising.[11]

Panel B repeats the analysis for the 3-year fundraising amount. The accelerator effect stands at 8.9%, and the 95% posterior interval is 7.4%–9.3%. The year effect is 0.0% (95% PI: 0.0%–0.2%), the industry effect is 1.5% (95% PI: 0.8%–2.6%). The pre-accelerator funding control explains 0.0% (95% PI: 0.0%–0.2%) of the variance in post-accelerator fundraising.[12]

The results are comparable to past variance-decomposition studies of startup performance. For example, Fitza et al. (2009) studied the startup valuation using a sample of US-based

[11]The estimate for the control variables' coefficient, β_s, is 0.1.

[12]The estimate for the control variables' coefficient, β_s, is 0.4.

startups. While not studying accelerator graduates per se, the study found an industry effect of 0.9%. Chan et al. (2020) explored the amount of funding raised within 1 year of entering an accelerator using a sample across the United States, Mexico, Kenya, and India (based on the Global Accelerator Learning Initiative (GALI) data). They report an industry effect of 0.0% and an accelerator effect of 11.1%.

Next, we proceed to estimate a full model where we add the manager and the cohort effects to the analysis. Panel A of Table 3 reports the Bayesian hierarchical estimates for fundraising amount within 12 months of entering an accelerator. We observe a reduction in the magnitude of the mean accelerator effect from 8.9% in the base model to 3.6% (95% PI: 2.4%–4.7%) in the full model. The manager effect stands at 7.7% (95% PI: 7.5%–7.9%). The cohort effect is of notable magnitude: 7.5% (95% PI: 7.4%–7.7%). The year effect remains at 0.0% (95% PI: 0.0%–0.6%). The industry effect is 3.2% (95% PI: 1.1%–4.6%), and pre-accelerator funding continues to explain 0.1% of the variance (95% PI: 0.0%–0.3%).

The results for fundraising amounts within 3 years of entering the accelerator are similar. As reported in Panel B, we find a reduction in the accelerator effect from 8.9% in the base model to 2.6% (95% PI: 1.1%–3.3%) in the full model. The manager effect stands at 4.9% (95% PI: 4.5%–5.4%), and the cohort effect is 5.6% (95% PI: 5.2%–6.0%). The industry effect is 1.6% (95% PI: 0.4%–2.3%), the year effect is 0.0% (95% PI: 0.0%–0.1%), and the control variable (pre-accelerator funding) remains at 0.0% of the variance (95% PI: 0.0%–0.3%).[13]

Finally, we conduct a set of additional analyses to check the robustness of the findings across different samples and methodologies (Appendix C). For example, we explored whether our results are driven by cohorts that have only one or a handful of startups that raised substantial funds. To that end, we follow Mangiola et al.'s (2021) approach to ascertain a Bayesian model's predictive abilities. Appendix C details this and other robustness tests and corresponding results.

Overall, our findings illustrate the magnitude of the cohort effect in comparison to the manager and accelerator effects. They suggest that the bottom-up mechanisms underlying the cohort effect are of similar (equally meaningful) in magnitude to the top-down manager effect and larger than the top-down accelerator effect. The next section interprets these results.

7 | DISCUSSION AND CONCLUSION

To paraphrase the title of this study, accelerators should not be viewed as solely akin to breweries; we document a significant vintage effect. In our primary analysis, we find that the cohort effect explains up to 7.5% of the variance in startup performance, while the accelerator effect explains between 2.6% and 3.6% once cohort and manager effects are included. We also find that the manager effect is associated with a notable magnitude of up to 7.7%. The posterior intervals for the effects are relatively small, suggesting that our sample size and model setup allow for accurate estimation of effect sizes.[14,15]

[13]The estimate for the controls' coefficient, β_s, is 0.0 for the 12-month model and 0.3% for the 3-year model.

[14]Neither the manager nor the cohort effect is within the 95% posterior intervals of the accelerator effect (and the accelerator effect is not within the 95% PI of manager and cohort effects), indicating that the accelerator effect is significantly different from the other two effects.

[15]In addition to the manager, accelerator, and cohort effect, our model also estimates the effect of industry, year, and the variance explained by the control variable of pre-accelerator funding. These last three effects are relatively small, suggesting they play a lesser role in explaining startup funding success.

The findings carry a few implications and contributions. First, we point to a new vehicle through which top-down mechanisms play out; specifically, the role of accelerator managers. We observe a reduction in the magnitude of the accelerator effect between the base model and the full model. Current discourse among scholars and practitioners tends to focus on the accelerator. The added value associated with each accelerator is often associated with the (top-down) design decisions it undertakes. Our findings call for a nuanced discourse. They suggest that managers play an important role in driving variation in the impact of top-down decisions.

Second, we find that cohorts account for a notable fraction of the variance in startup performance. We recognize that business and scholarly discussions often celebrate the accelerators and attribute startups' success to the training and networking that accelerators offer. Our study indicates that one should also acknowledge significant bottom-up cohort dynamics and their role in the success of accelerators' participants. In the spirit of variance-decomposition studies, we direct scholars to potential fertile areas for further investigation. We underscore the need to expand on recent cohort dynamics studies, such as Assenova and Amit (2024), Cohen, Bingham, and Hallen (2019), Dushnitsky and Sarkar (2018), and Hallen et al. (2020). Doing so will advance our understanding of when, why and how some cohorts turn into thriving communities. Along these lines, Table 1 offers guidance on a plethora of plausible mechanisms. A corollary observation for practitioners is to monitor cohort dynamics because it accounts for a notable fraction of performance.

These theoretical contributions are possible through a parallel effort in data and methodology. The data extend prior work by looking beyond startups' accelerator-affiliation, to further capture the effect of managers and cohorts on startups' performance. Methodologically, we present an approach to variance-decomposition analysis that is new to the strategy field to the best of our knowledge. It joins ANOVA and maximum likelihood estimation (MLM) approaches that informed past studies of established firms performance (e.g., McGahan & Porter, 1997; Misangyi et al., 2006; Quigley & Graffin, 2017; Vanneste, 2017). We estimate a Bayesian hierarchical model (Gelman et al., 2013) using a Tweedie distribution to model our highly right-skewed zero-inflated dependent variable. The study details the setup, estimation, and presentation of results which can be helpful for other scholars studying non-normally distributed performance outcomes or those interested in settings characterized by sparseness in hierarchical structures.

A few notable caveats are warranted. As with all variance-decomposition analyses, our study only captures effect sizes and not their underlying causes (McGahan & Porter, 1997; McGahan & Victer, 2010). We view this less as a limitation and more as a feature of variance-decomposition work. It serves as a "call to action." Namely, this type of work advances research by highlighting the relative importance of effect classes, and it does so by bundling together a multitude of factors that might interact in causally complex ways (Vanneste, 2017; Vedula & Fitza, 2019). Another feature of variance-decomposition work relates to the approach to thinking about selection. This is often tackled through careful interpretation; that is, variance-decomposition studies focus on estimating whether or not the performance of certain firm populations differs, conditional on initial selection decisions (Karniouchina et al., 2013; Vedula & Fitza, 2019).[16] In the context of our accelerator study, the selection is fully attributed to top-down mechanisms because it is the accelerator and their managers that screen and attract prospective startups. In other words, the cohort-effect estimates can be

[16]To see this, consider the fact that past variance-decomposition works estimate industry- and business-unit effects based on the assumption that corporations choose which industries to compete in and which business units to develop or acquire. The interpretation of an estimate of the corporate effect, therefore, is that it captures the fraction of profitability variation attributed to a corporation's ability to select as well as operate the business unit.

interpreted as directly capturing the bottom-up mechanisms associated with the dynamics among the tightly-knit cohort members.

DATA AVAILABILITY STATEMENT

Restrictions apply to the availability of these data, which were used under license for this study.

ORCID

Gil Avnimelech https://orcid.org/0000-0002-7102-809X
Markus Fitza https://orcid.org/0000-0002-9422-7032

REFERENCES

Assenova, V. A. (2020). Early-stage venture incubation and mentoring promote learning, scaling, and profitability among disadvantaged entrepreneurs. *Organization Science*, *31*(6), 1560–1578.

Assenova, V. A., & Amit, R. (2024). Poised for growth: Exploring the relationship between accelerator program design and startup performance. *Strategic Management Journal*, *45*(6), 1029–1060.

Avnimelech, G., & Rechter, E. (2023). How and why accelerators enhance female entrepreneurship. *Research Policy*, *52*(2), 104669.

Avnimelech, G., & Teubal, M. (2006). Creating venture capital industries that co-evolve with high tech: Insights from an extended industry life cycle perspective of the Israeli experience. *Research Policy*, *35*(10), 1477–1498.

Banc, C., & Messeghem, K. (2020). Discovering the entrepreneurial micro-ecosystem: The case of a corporate accelerator. *Thunderbird International Business Review*, *62*(5), 593–605.

Beaumont, M. A., & Rannala, B. (2004). The Bayesian revolution in genetics. *Nature Reviews Genetics*, *5*(4), 251–261.

Box, G. E. P., & Draper, N. R. (1987). *Empirical model-building and response surface*. Wiley.

Bueckner, P. (2024). Prior Definitions for brms Models. https://paul-buerkner.github.io/brms/reference/set_prior.html

Bischoff, K. M., Gielnik, M. M., & Frese, M. (2020). When capital does not matter: How entrepreneurship training buffers the negative effect of capital constraints on business creation. *Strategic Entrepreneurship Journal*, *14*(3), 369–395.

Cai, J., & Szeidl, A. (2018). Interfirm relationships and business performance. *The Quarterly Journal of Economics*, *133*(3), 1229–1282.

Camuffo, A., Cordova, A., Gambardella, A., & Spina, C. (2020). A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science*, *66*(2), 564–586.

Chan, C. S. R., Patel, P. C., & Phan, P. H. (2020). Do differences among accelerators explain differences in the performance of member ventures? Evidence from 117 accelerators in 22 countries. *Strategic Entrepreneurship Journal*, *14*(2), 224–239.

Chatterji, A., Delecourt, S., Hasan, S., & Koning, R. (2019). When does advice impact start-up performance? *Strategic Management Journal*, *40*(3), 331–356.

Clarysse, B., & Yusubova, A. (2014). Success factors of business accelerators. In *Technology business incubation mechanisms and sustainable regional development, proceedings*. Toulouse Business School.

Clarysse, B., Wright, M., & Van Hove, J. (2015). *A look inside accelerators*. Nesta.

Clingingsmith, D., & Shane, S. (2018). Training aspiring entrepreneurs to pitch experienced investors: Evidence from a field experiment in the United States. *Management Science*, *64*(11), 5164–5179.

Clough, D. R., Fang, T. P., Vissa, B., & Wu, A. (2019). Turning lead into gold: How do entrepreneurs mobilize resources to exploit opportunities? *Academy of Management Annals*, *13*(1), 240–271.

Cohen, S. L., & Hochberg, Y. V. (2014). *Accelerating start-ups: The seed accelerator phenomenon*. *SSRN Journal*, 1–16. https://doi.org/10.2139/ssrn.2418000

Cohen, S. L., Bingham, C. B., & Hallen, B. L. (2019). The role of accelerator designs in mitigating bounded rationality in new ventures. *Administrative Science Quarterly*, *64*(4), 810–854.

Cohen, S. L., Fehder, D. C., Hochberg, Y. V., & Murray, F. (2019). The design of start-up accelerators. *Research Policy*, *48*(7), 1781–1797.

Conti, A. (2018). Entrepreneurial finance and the effects of restrictions on government R&D subsidies. *Organization Science*, *29*(1), 134–153.

Conti, R., Kacperczyk, O., & Valentini, G. (2022). Institutional protection of minority employees and entrepreneurship: Evidence from the LGBT employment non-discrimination acts. *Strategic Management Journal*, *43*(4), 758–791.

Crawford, G. C., Aguinis, H., Lichtenstein, B., Davidsson, P., & McKelvey, B. (2015). Power law distributions in entrepreneurship: Implications for theory and research. *Journal of Business Venturing*, *30*(5), 696–713.

Crişan, E. L., Salanță, I. I., Beleiu, I. N., Bordean, O. N., & Bunduchi, R. (2021). A systematic literature review on accelerators. *The Journal of Technology Transfer*, *46*(1), 62–89.

Drori, I., & Wright, M. (2018). Accelerators: Characteristics, trends and the new entrepreneurial ecosystem. In M. Wright & I. Drori (Eds.), *Accelerators*. Edward Elgar Publishing.

Dunn, P. K., & Smyth, G. K. (1996). Randomized quantile residuals. *Journal of Computational and graphical statistics*, *5*(3), 236–244.

Dunn, P. K., & Smyth, G. K. (2001). Tweedie family densities: Methods of evaluation. In *Proceedings of the 16th international workshop on statistical modelling* (pp. 2–6).

Dushnitsky, G., & Sarkar, S. (2018). Variance decomposing of accelerator and cohort effects among London startups. *Academy of Management Best Papers Proceedings*, *2018*, 10522.

Dushnitsky, G., & Sarkar, S. (2022). Here comes the sun: The impact of incidental contextual factors on entrepreneurial resource acquisition. *Academy of Management Journal*, *65*(1), 66–92.

El-Shaarawi, A. H., Zhu, R., & Joe, H. (2011). Modelling species abundance using the Poisson–Tweedie family. *Environmetrics*, *22*(2), 152–164.

Engel, J. S., & del-Palacio, I. (2011). Global clusters of innovation: The case of Israel and Silicon Valley. *California Management Review*, *53*(2), 27–49.

Fitza, M. A. (2014). The use of variance decomposition in the investigation of CEO effects: How large must the CEO effect be to rule out chance? *Strategic Management Journal*, *35*(12), 1839–1852.

Fitza, M. A., Matusik, S. F., & Mosakowski, E. (2009). Do VCs matter? The importance of owners on performance variance in start-up firms. *Strategic Management Journal*, *30*(4), 387–404.

Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., & Gelman, A. (2019). Visualization in Bayesian workflow. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *182*(2), 389–402.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). 515-534.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian data analysis* (2nd ed.). Chapman and Hall.

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical science*, *7*(4), 457–472.

Gelman, A., & Pardoe, I. (2006). Bayesian measures of explained variance and pooling in multilevel (hierarchical) models. *Technometrics*, *48*(2), 241–251.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). CRC Press.

Gilchrist, R., & Drinkwater, D. (2000). The use of the Tweedie distribution in statistical modelling. In *COMPSTAT: Proceedings in computational statistics 14th symposium held in Utrecht, the Netherlands, 2000* (pp. 313–318). Physica-Verlag HD.

Glinik, M. (2019). Gruendungsgarage—A best-practice example of an academic start-up accelerator. *International Journal of Engineering Pedagogy*, *9*(3), 33–43.

González-Uribe, J., & Leatherbee, M. (2018). The effects of business accelerators on venture performance: Evidence from start-up Chile. *Review of Financial Studies*, *31*(4), 1566–1603.

González-Uribe, J., & Reyes, S. (2021). Identifying and boosting "gazelles": Evidence from business accelerators. *Journal of Financial Economics*, *139*(1), 260–287.

Guler, I. (2007). Throwing good money after bad? Political and institutional influences on sequential decision making in the venture capital industry. *Administrative Science Quarterly*, *52*(2), 248–285.

Guzman, J., & Stern, S. (2016). *The state of American entrepreneurship: New estimates of the quantity and quality of entrepreneurship for 15 US states, 1988–2014* (NBER Working Paper No. 22095). National Bureau of Economic Research.

Hallen, B. L. (2008). The causes and consequences of the initial network positions of new organizations: From whom do entrepreneurs receive investments? *Administrative Science Quarterly*, *53*(4), 685–718.
Hallen, B. L., Cohen, S. L., & Bingham, C. B. (2020). Do accelerators work? If so, how? *Organization Science*, *31*(2), 378–414.
Hambrick, D. C., & Quigley, T. J. (2014). Toward more accurate contextualization of the CEO effect on firm performance. *Strategic Management Journal*, *35*(4), 473–491.
Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, *15*(1), 1593–1623.
Howell, T. (2022). Coworking spaces: An overview and research agenda. *Research Policy*, *51*(2), 104447.
Huang, Y., Shen, C., & Contractor, N. S. (2013). Distance matters: Exploring proximity and homophily in virtual world networks. *Decision Support Systems*, *55*(4), 969–977.
Jorgensen, B. (1987). Exponential dispersion models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *49*(2), 127–162.
Jorgensen, B. (1997). *The theory of dispersion models*. Chapman & Hall. CRC Monographs on Statistics and Applied Probability.
Karniouchina, E. V., Carson, S. J., Short, J. C., & Ketchen, D. J., Jr. (2013). Extending the firm vs. industry debate: Does industry life cycle stage matter? *Strategic Management Journal*, *34*(8), 1010–1018.
Kohler, T. (2016). Corporate accelerators: Building bridges between corporations and startups. *Business Horizons*, *59*(3), 347–357.
Kokonendji, C. C., Bonat, W. H., & Abid, R. (2021). Tweedie regression models and its geometric sums for (semi-) continuous data. *Wiley Interdisciplinary Reviews: Computational Statistics*, *13*(1), e1496.
Korteweg, A., & Sorensen, M. (2017). Skill and luck in private equity performance. *Journal of Financial Economics*, *124*(3), 535–562.
Krishnan, R., Cook, K. S., Kozhikode, R. K., & Schilke, O. (2021). An interaction ritual theory of social resource exchange: Evidence from a silicon valley accelerator. *Administrative Science Quarterly*, *66*(3), 659–710.
Lall, S., Bowles, L., & Ross, B. (2013). Bridging the "Pioneer Gap": The role of accelerators in launching high-impact enterprises. *Innovations: Technology, Governance, Globalization*, *8*(3), 105–137.
Leatherbee, M., & Katila, R. (2020). The lean start-up method: Early-stage teams and hypothesis-based probing of business ideas. *Strategic Entrepreneurship Journal*, *14*(4), 570–593.
Lerner, J. (1994). The syndication of venture capital investments. *Financial Management*, *23*, 16–27.
Mangiola, S., Thomas, E. A., Modrák, M., Vehtari, A., & Papenfuss, A. T. (2021). Probabilistic outlier identification for RNA sequencing generalized linear models. *NAR Genomics and Bioinformatics*, *3*(1), lqab005.
Mansoori, Y., Karlsson, T., & Lundqvist, M. (2019). The influence of the lean startup methodology on entrepreneur–coach relationships in the context of a startup accelerator. *Technovation*, *84*, 37–47.
Matusik, S. F., & Fitza, M. A. (2012). Diversification in the venture capital industry: Leveraging knowledge under uncertainty. *Strategic Management Journal*, *33*, 407–426.
McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in R and Stan*. Chapman and Hall/CRC.
McGahan, A. M., & Porter, M. E. (1997). How much does industry matter, really? *Strategic Management Journal*, *18*(S1), 15–30.
McGahan, A. M., & Victer, R. (2010). How much does home country matter to corporate profitability? *Journal of International Business Studies*, *41*(1), 142–165.
Mikey, T. (2016). *One-third of U.S. startups that raised a Series A in 2015 went through an accelerator*. Pitchbook. https://pitchbook.com/news/articles/one-third-of-us-startups-that-raised-a-series-a-in-2015-went-through-an-accelerator
Miller, P., & Bound, K. (2011). *The startup factories*. NESTA. http://www.nesta.org.uk/library/documents/StartupFactories.pdf
Misangyi, V. F., Elms, H., Greckhamer, T., & Lepine, J. A. (2006). A new perspective on a fundamental debate: A multilevel approach to industry, corporate, and business unit effects. *Strategic Management Journal*, *27*(6), 571–590.
Mullahy, J. (1986). Specification and testing of some modified count data models. *Journal of Econometrics*, *33*(3), 341–365.

Nickerson, J. A., & Zenger, T. R. (2008). Envy, comparison costs, and the economic theory of the firm. *Strategic Management Journal*, *29*(13), 1429–1449.

Paez, E., Faraway, J. J., & Davison, A. C. (2008). Bayesian hierarchical models with an application to air pollution measurement. *Environmetrics*, *19*(7), 665–682.

Prexl, K., Hubert, M., Beck, S., Heiden, C., & Prügl, R. (2019). Identifying and analysing the drivers of heterogeneity among ecosystem builder accelerators. *R&D Management*, *49*(4), 624–638.

Probst, G., & Borzillo, S. (2008). Why communities of practice succeed and why they fail. *European Management Journal*, *26*(5), 335–347.

Quigley, T. J., & Graffin, S. D. (2017). Reaffirming the CEO effect is significant and much larger than chance: A comment on Fitza (2014). *Strategic Management Journal*, *38*(3), 793–801.

Radojevich-Kelley, N., & Hoffman, D. L. (2012). Analysis of accelerator companies: An exploratory case study of their programs, processes, and early results. *Small Business Institute Journal*, *8*(2), 54–70.

Rechter, E., & Avnimelech, G. (2024). Intensive personal mentoring: Accelerators' secret sauce. *Small Business Economics*, 1–26.

Roche, M., Ottel, A., & Catalini, C. (2022). *(Co-)working in close proximity: Knowledge spillovers and social interactions*. HBS Working Paper 21-024.

Scherer, F. M., Harhoff, D., & Kukies, J. (2000). Uncertainty and the size distribution of rewards from innovation. *Journal of Evolutionary Economics*, *10*(1–2), 175–200.

Schmalensee, R. (1985). Do markets differ much? *American Economic Review*, *75*(3), 341–351.

Senor, D., & Singer, S. (2011). *Start-up nation: The story of Israel's economic miracle*. McClelland & Stewart.

Shankar, R. K., & Clausen, T. H. (2020). Scale quickly or fail fast: An inductive study of acceleration. *Technovation*, *98*, 102174.

Sharma, A., Joshi, M., & Shukla, B. (2014). *Is accelerator an option? Impact of accelerator in start-up eco-system!* Available at SSRN: https://ssrn.com/abstract=2438846

StartupBlink. (2021). *The Global Startup Ecosystem Index Report 2021*. https://report.startupblink.com/

Stavropoulos, S., Burger, M. J., & Skuras, D. (2015). Data sparseness and variance in accounting profitability. *Organizational Research Methods*, *18*(4), 656–678.

Stuart, T. E., Hoang, H., & Hybels, R. C. (1999). Interorganizational endorsements and the performance of entrepreneurial ventures. *Administrative Science Quarterly*, *44*(2), 315–349.

Swallow, B., Buckland, S. T., King, R., & Toms, M. P. (2016). Bayesian hierarchical modelling of continuous non-negative longitudinal data with a spike at zero: An application to a study of birds visiting gardens in winter. *Biometrical Journal*, *58*(2), 357–371.

Thompson, M. (2005). Structural and epistemic parameters in communities of practice. *Organization Science*, *16*(2), 151–164.

Tweedie, M. C. K. (1984). An index which distinguishes between some important exponential families. In J. K. Ghosh & J. Roy (Eds.), *Statistics: Applications and new directions. Proceedings of the Indian Statistical Institute Golden Jubilee International Conference* (pp. 579–604). Indian Statistical Institute.

Uhm, C. H., Sung, C. S., & Park, J. Y. (2018). Understanding the accelerator from resources-based perspective. *Asia Pacific Journal of Innovation and Entrepreneurship*, *12*(3), 258–278.

Vanneste, B. S. (2017). How much do industry, corporation, and business matter, really? A meta-analysis. *Strategy Science*, *2*(2), 121–139.

Vedula, S., & Fitza, M. (2019). Regional recipes: A configurational analysis of the regional entrepreneurial ecosystem for US venture capital-backed startups. *Strategy Science*, *4*(1), 4–24.

Vehtari, A. (2019). *Why Student T(3,0,1) for prior*? https://discourse.mc-stan.org/t/why-studentt-3-0-1-for-prior/8102

Vehtari, A., Simpson, D., Gelman, A., Yao, Y., & Gabry, J. (2015). Pareto smoothed importance sampling. *arXiv preprint arXiv:1507.02646*.

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*, 1413–1432.

Wasserman, L. (2004). *All of nonparametric statistics*. Springer Science & Business Media.

Wenger, E. C. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.

Wenger, E. C., McDermott, R., & Snyder, W. M. (2002). *Cultivating communities of practice*. Harvard Business School Press.

Wibbens, P. D. (2019). Performance persistence in the presence of higher-order resources. *Strategic Management Journal*, *40*(2), 181–202.

Withers, M. C., & Fitza, M. A. (2017). Do board chairs matter? The influence of board chairs on firm performance. *Strategic Management Journal*, *38*(6), 1343–1355.

Wright, M., Siegel, D. S., & Mustar, P. (2017). An emerging ecosystem for student start-ups. *The Journal of Technology Transfer*, *42*(4), 909–922.

Yang, S., Kher, R., & Lyons, T. S. (2018). Where do accelerators fit in the venture creation pipeline? Different values brought by different types of accelerators. *Entrepreneurship Research Journal*, *8*(4), 70–140.

Ye, T., Lachos, V. H., Wang, X., & Dey, D. K. (2021). Comparisons of zero-augmented continuous regression models from a Bayesian perspective. *Statistics in Medicine*, *40*(5), 1073–1100.

Yu, S. (2020). How do accelerators impact the performance of high-technology ventures? *Management Science*, *66*(2), 530–552.

Zhang, Y., & Liu, Y. (2019). A comparison of maximum likelihood and Bayesian estimation methods for zero-inflated Poisson regression models. *Journal of Statistical Computation and Simulation*, *89*(3), 445–461.

How to cite this article: Avnimelech, G., Dushnitsky, G., Ellsaesser, F., & Fitza, M. (2024). Are accelerators akin to breweries or wineries? A Bayesian variance decomposition of accelerator and cohort effects. *Strategic Management Journal*, 1–46. https://doi.org/10.1002/smj.3665

APPENDIX A: BAYESIAN MODEL SPECIFICATION: PRIORS AND STARTING VALUES

As described in Equation (1) in the main text, we use a Tweedie distribution for the dependent variable:

$$y_s \sim Tw(\mu_s, \varphi, p). \tag{A.1}$$

The parameters of the Tweedie distribution are the mean μ_s dispersion φ and power parameter p (in the STAN code it appears as mtheta (McElreath, 2020), see Appendix D), where each of the parameters has to take non-negative values (El-Shaarawi et al., 2011; Gilchrist & Drinkwater, 2000; Kokonendji et al., 2021; Ye et al., 2021). $\varphi \sim gamma(k=0.01, \theta=0.01)$: The gamma distribution here ensures positive values in a realistic range for φ (Gelman et al., 2013). p is constrained to be in the interval $1 < p < 2$, as this ensures a zero inflated distribution (El-Shaarawi et al., 2011; Ye et al., 2021) (adding this constraint increases sampling efficiency). The variance of the Tweedie distribution is defined as (Swallow et al., 2016):

$$Var(y_s) = \varphi * \mu_s^p.$$

The estimation of μ_s is based on Equation (2) in the paper. Following Gelman (2006), for the variances of the random effects (σ_γ, $\sigma_{\delta,a}$, $\sigma_{\varphi,m}$, and σ_i), we use half-t distributed hyper-priors, with: $Student\ half - T(\nu, \mu_T, \sigma_T)$ on the range [0, inf), where ν stands for the degrees of freedom, μ_T for the mean, and σ_T for the standard deviation of the t-distribution. The half-t distributed hyper-priors have been chosen as some of the levels only have a few observations, and an unconstrained (flat) hyper-prior would thus not result in sufficient shrinkage for these levels (Gelman, 2006, p. 528 ff). Per prior work (Bueckner, 2024; Vehtari, 2019), we set ν at 3 because having degrees of freedom greater than 2 ensures a finite variance and mean while for values larger than 3 the tails become extremely heavy and difficult to bound. The mean, μ_T, is set at zero and σ_T is set at 2.5 (Bueckner, 2024; Vehtari, 2019). We tested sensitivities to varying the degrees of freedom of ν to 2 and of σ_T to 1, which does not affect our results. Per Gelman et al. (2013), the intercept, β_0, year fixed effect, $\beta_{y(s)}$, and previous funding, β_s, have a flat prior on (inf, inf).

The results are based on posterior inference of 2 Markov chains sampled with the No U Turn (NUTS) Hamiltonian Monte Carlo sampler (Hoffman & Gelman, 2014). For each chain 4000 draws are performed, of which the first 2000 are taken as warm-up. No thinning is performed as the samples can be kept in memory. This results in a total of (2 times 2000) 4000 posterior draws for the model. We also run a setup with 4 Markov chains (2000 iterations for warm-up, and 2000 iterations for sampling) for the main model and it did not change the results.

The variance explained, and the confidence intervals (called posterior intervals in the context of our Bayesian model) of the random intercepts can be directly taken from the Bayesian posteriors. The variance explained by the fixed year effect and by the control variable is calculated by once taking the model with and once without the associated variables and calculating the difference in the explained variance (Misangyi et al., 2006; Quigley & Graffin, 2017; Withers & Fitza, 2017)[17].

[17] The posterior intervals for these two effects are calculated via bootstrapping with 100 iterations.

The variance explained by a component (i.e., our effect of interest, such as the accelerator effect) is defined as:

$$\frac{\sigma^2_{component}}{total\, variance}.$$

For $\sigma^2_{component}$ we use the mean across all posterior draws. The *total variance* is the sum of the individual variances plus the error term. Recall that an advantage of the Bayesian analysis is the ability to accommodate varying variances for the manager and the cohort (e.g., allowing the variance of the manager to differ across accelerators, and similarly allowing the variance of the cohort to differ across managers). Therefore, the $\sigma^2_{component}$ for the manager-effect is based on the mean of the set of variance values estimated across the different accelerators, and the $\sigma^2_{component}$ for the cohort effect is based on the mean variance values estimated across the different managers.

APPENDIX B: CONVERSION AND VALIDATION OF THE BAYESIAN MODEL

We undertake several steps to assess whether the model fits with the data. To that end, we follow the suggestions of Gelman et al. (2004, p. 281), Wibbens (2019), Vehtari et al. (2017), and Vehtari et al. (2015) regarding which post-estimation tests to conduct.

Effective sample size

In Bayesian statistics, the effective sample size is a measure of the amount of information contained in a sample. It represents the size of an equivalent independent and identically distributed (IID) sample that would provide the same amount of information as the actual sample (Wasserman, 2004). For each parameter of interest, n_{eff} should be at least 10 times the number of Markov chains (see Gelman et al., 2013, p. 287; Wibbens, 2019). In our case, the effective sample size is between 842 and 1839, which is more than the required 20. We report our effective sample sizes in Table B1.

TABLE B1 Key convergence statistics for the models.

	DV: Funding raised within 12-months	DV: Funding raised within 3-years
$\hat{R}$ Sector	1.001	1.001
$\hat{R}$ Accelerator	1.002	1.004
$\hat{R}$ Manager	1.001	1.002
$\hat{R}$ Cohort	1.001	1.004
Sector_{eff}	1839	1759
Accelerator_{eff}	981	842
Manager_{eff}	1287	1285
Cohort_{eff}	1215	1288

Convergence of the Markov chains

In the field of Bayesian statistics, assessing the convergence of the Markov chains is essential for an accurate analysis. $\hat{R}$, or the potential scale reduction factor, serves as a diagnostic tool for this purpose. This quantitative measure, developed by Gelman and Rubin (1992), compares the within-chain and between-chain variability of multiple chains in Markov chain Monte Carlo (MCMC) simulations, each initialized with different starting values. It is important to ensure that the $\hat{R}$ statistic is sufficiently low to have confidence in the accuracy of the posterior distribution obtained from the MCMC simulation (Gelman et al., 2013, p. 287). The $\hat{R}$ statistic needs to be below 1.10 for all variables to ensure that the Markov chains have mixed well and that the MCMC algorithm is not exploring an incorrect or incomplete posterior distribution (Gelman et al., 2013). In all our models $\hat{R}$ for the variables sector, accelerator, manager, and cohort ranges between 1.001 and 1.004 and is below the 1.10 threshold, suggesting that the Markov chains converged and mixed well (see Table B1).

The trajectory of the Markov chains

The models' convergence can also be evaluated visually by examining the trajectory of the Markov chains. During the warm-up phase, the chains' values may be non-stationary and change over time, but they should stabilize and become stationary after this phase (Gelman et al., 2013). Additionally, the mixing of the Markov chains can be assessed visually, as they should oscillate around similar values, ensuring that each sampled trajectory does not yield different outcomes (see Gelman et al., 2013, p. 283, fig. 11.3). Figure B1 shows two Markov chains (blue and black), each sampled 2000 times after warm-up. The two Markov chains are stationary and oscillate in the same parameter range, which means that the chains are mixing as required (Gelman et al., 2013, p. 283), confirming the $\hat{R}$ statistics provided above. Note that the chains are truncated at zero, as positive values are enforced by the Tweedie distribution.

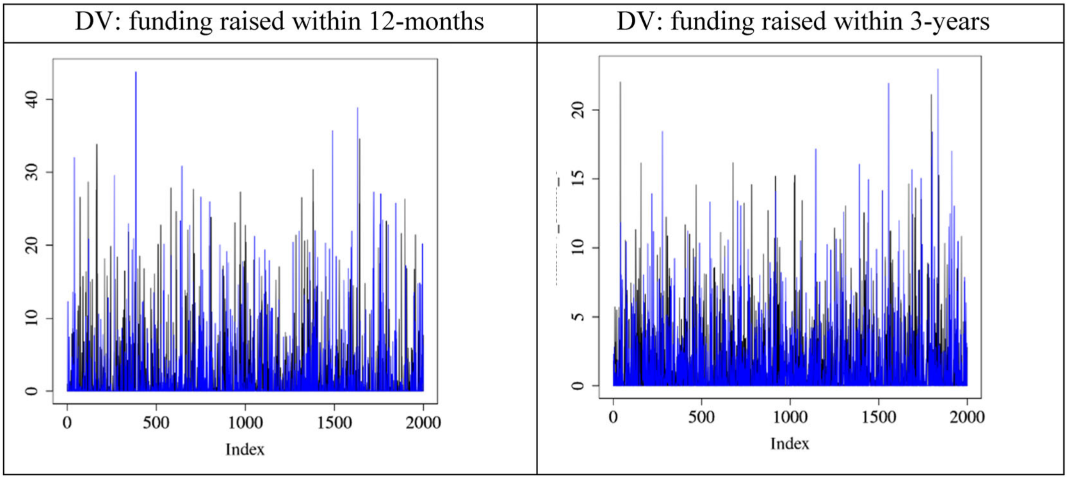

FIGURE B1 Two Markov chains (black and blue) for the 12-month and 3-year DV after the warm-up period. Black line: Markov Chain 1; Blue line: Markov Chain 2.

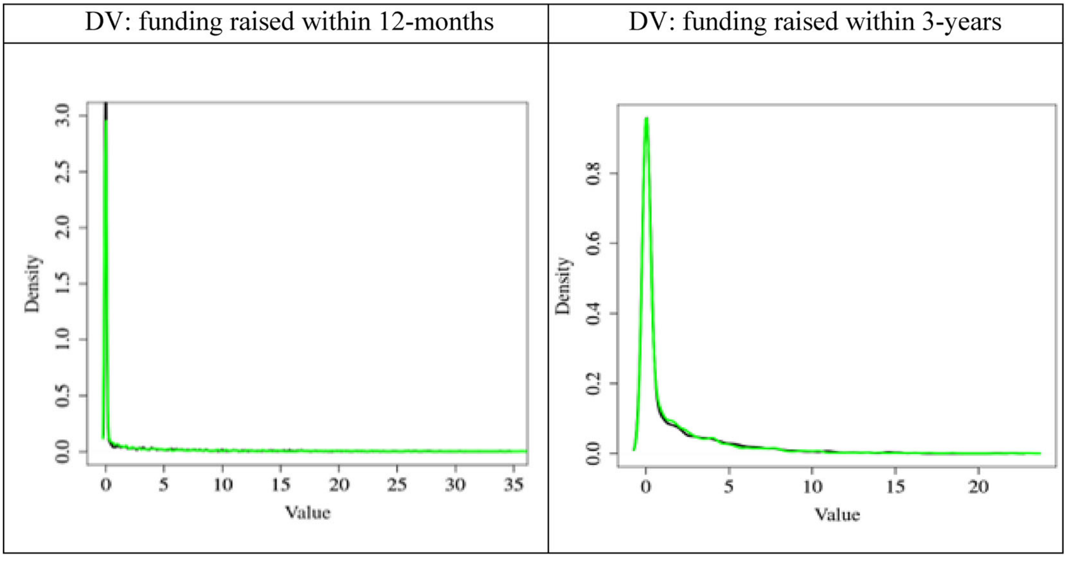

FIGURE B2 Density of Markov chain for 12-months and 3-years after the warm-up period. Green line: Markov Chain 1; Blue line: Markov Chain 2.

The sampling after the warm-up for the Markov chains can also be represented as a density plot, which is shown in Figure B2. The two Markov chains are plotted in blue and black and closely lie on top of each other. This confirms that the Markov chains are mixing, as indicated in the previous figures. Taken together, the Markov chains, $\hat{R}$, and effective sample size indicate that the model has converged well.

Distribution of the error term

Next, we assess the distribution of the error term (Gelman et al., 2013). One effective way to evaluate the fit of the error term distribution is through the use of Bayesian residuals (Dunn & Smyth, 1996). A distribution of the residuals that is symmetric and centered at zero is an indication of a good model fit because it suggests that the model's assumptions are consistent with the observed data and that the error term distribution is properly specified (Box & Draper, 1987). Figure B3 plots the error distribution terms of the 12-month and 3-year DV models (pooled across all posterior simulations) (Gelman et al., 2013). For the 12-month DV model $n = 4000$ draws * 1350 observation (5.4 million). For the 3-year DV model $n = 4000$ draws * 515 observation (2.06 million). The error is zero centered and symmetrically distributed indicating that the model represents the dependent variable well.

Distribution generated by the posteriors

Additionally, we can compare the statistics of the distributions of the DV generated by the posteriors of the two fit models to the distribution of the DV of the original data (Gelman et al., 2013). Figure B4 shows the means of the 12-month DV, and the 3-year DV of the original

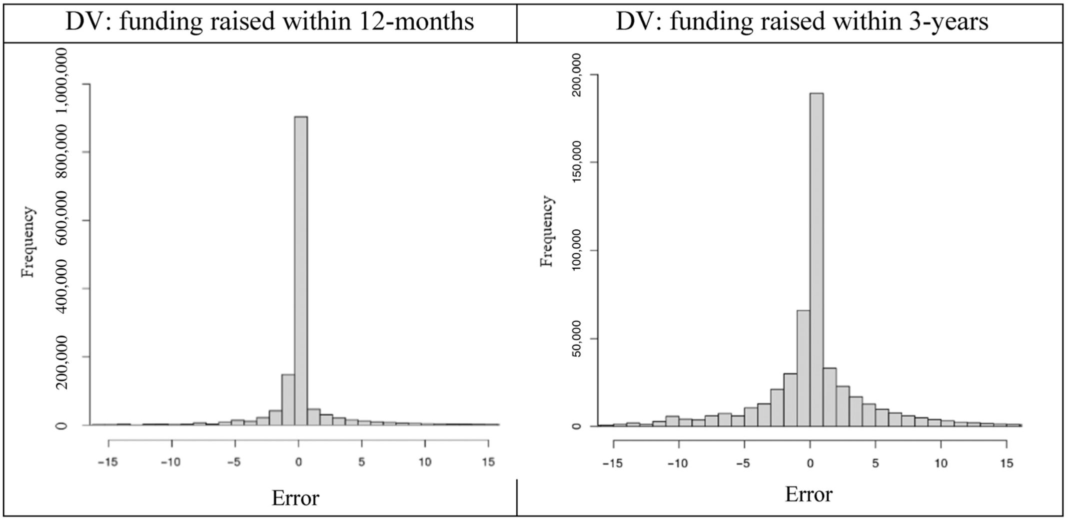

FIGURE B3 Distribution of the error term.

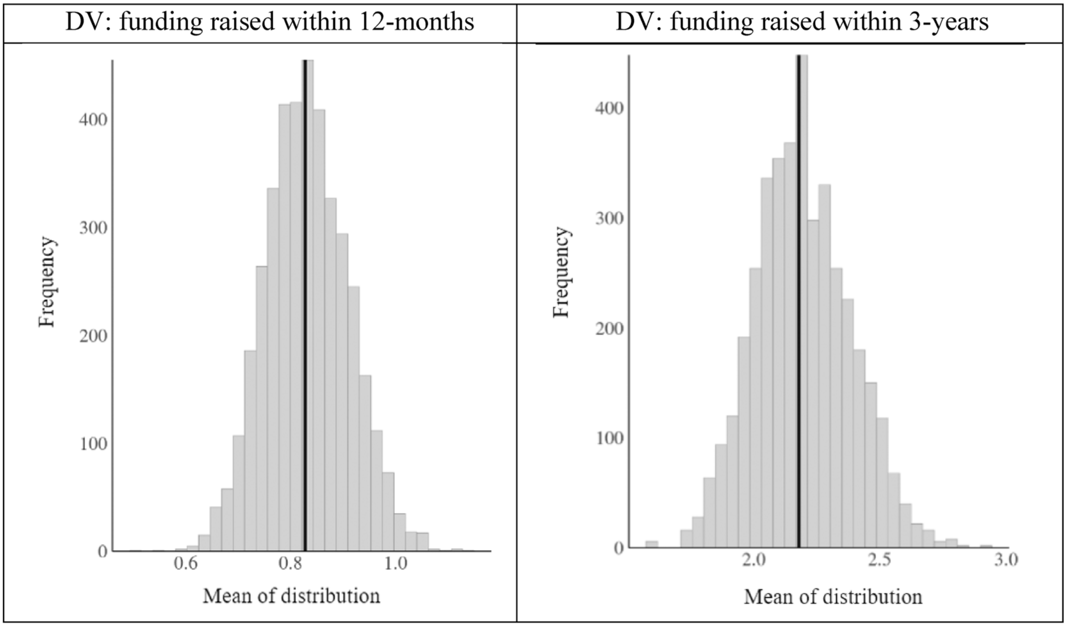

FIGURE B4 Mean of the posterior compared to mean of data. Solid black line: Mean of the DV in the data. Gray shaded bars: Distribution of the means of 4000 simulations of each model.

data as solid black horizontal lines in comparison to the means of 4000 simulations of each of the models (gray shaded bars). The distribution of the means of the posterior draws of the models has peaks around the means of the DV of the data and is symmetrically distributed, indicating that the Bayesian models estimate the means of the DVs well.

Extreme value statistics

Another way to validate the model is to check extreme value statistics, the maximum and minimum values of draws of the posterior compared to the maximum and minimum values of the dependent variable (Gelman et al., 2013).

The *minimum value* is zero for the dependent variable. Thus Figure B5 compares the number of zeros for the DVs with the number of zeros in the 4000 posterior model draws.

The number of zeros for the 12-month DV is 1020 (solid black line). The number of zeros for the 3-year DV is 249 (solid black line). The posterior draws (gray shaded bars) cluster tightly around these respective values, indicating that the models represent the zero inflation well.

Figure B6 shows the *maximum values* of 4000 model draws of the posterior distributions compared to the maximum values of the 12-month DV, which is 60, and the maximum value of the 3-year DV, which is 109. In each case the model represents the extreme values well. The slight skew of the extreme value distribution to the right can be expected, as the distribution generated by the Tweedie model is right-skewed (El-Shaarawi et al., 2011, Gilchrist & Drinkwater, 2000, Kokonendji et al., 2021, Ye et al., 2021).

Density overlay plot

The fit of the posterior of the models with the dependent variables (i.e., fundraising 12-months and 3-years after joining an accelerator) is confirmed by a density overlay plot presented in Figure B7 (Gelman et al., 2013, chap 6; Gabry et al., 2019). The dark red line represents the

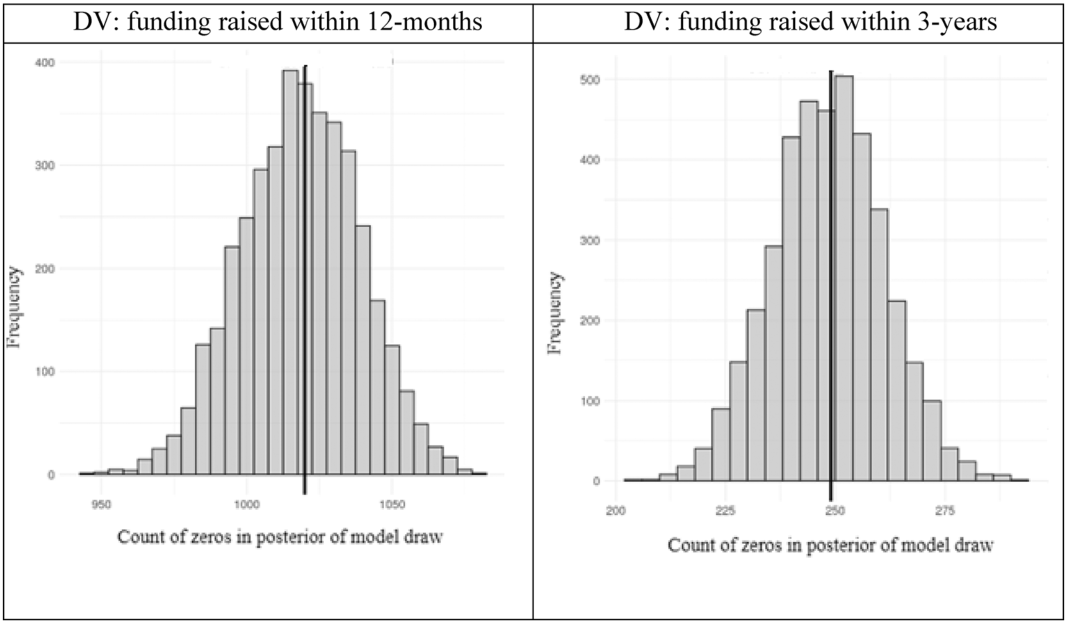

FIGURE B5 Number of zeros generated by model draws versus original data. Solid black line: Number of zeros for the DV in the data. Gray shaded bars: Distribution of how many zeros there are in the posterior draws of 4000 simulations of each model.

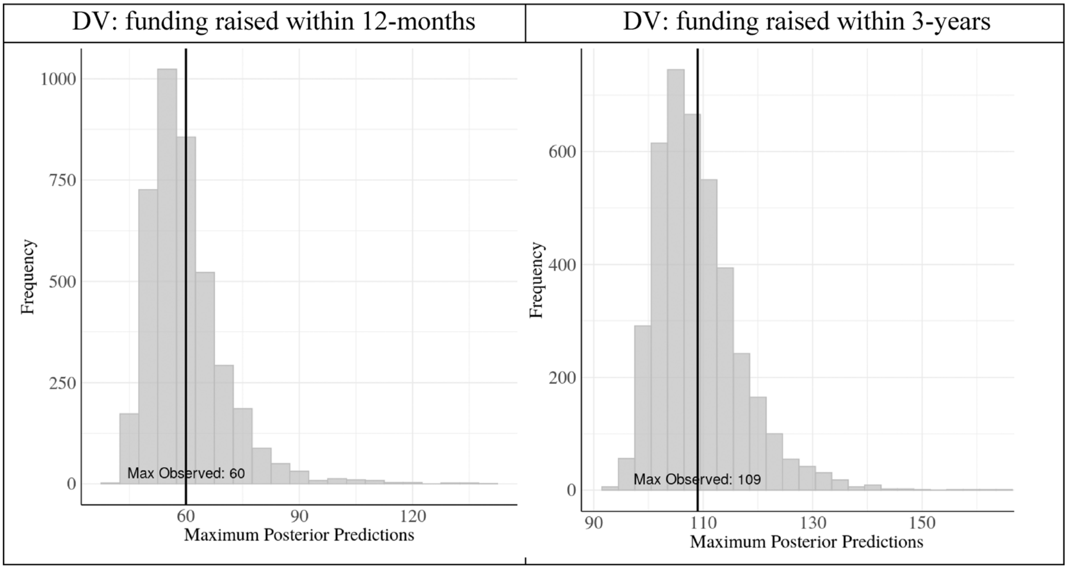

FIGURE B6 Maximum value of the different model draws versus original data. Solid black line: Maximum value observed in data. Gray shaded bars: Distribution of maximum values in the posterior draws of 4000 simulations of each model.

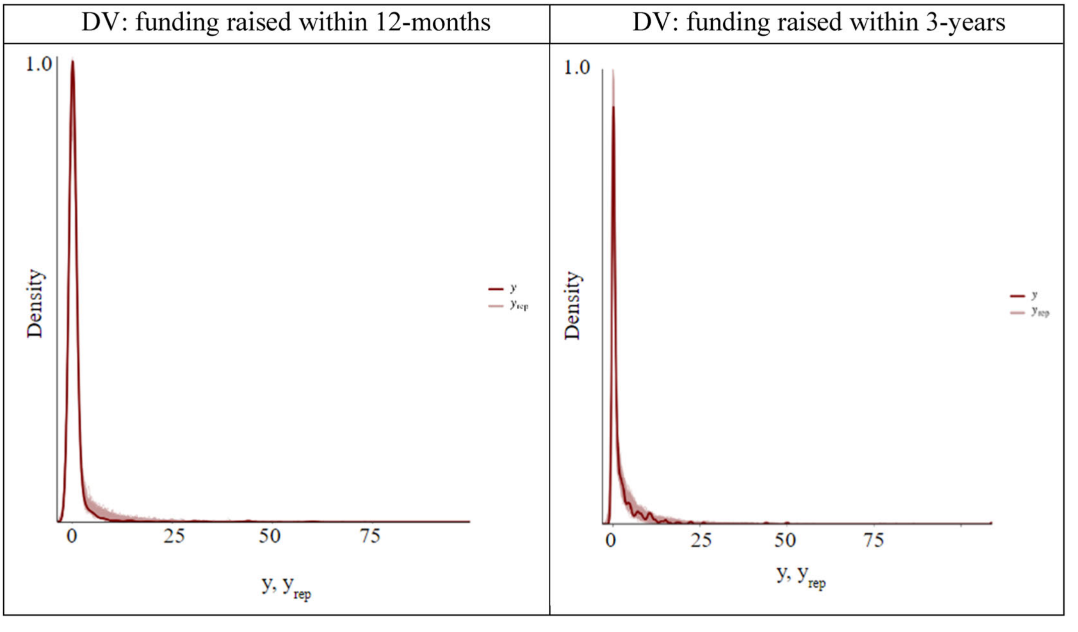

FIGURE B7 Density overlay of DVs compared to posterior draws. Dark red line: Density of the dependent variable. Red shaded area: Density of the distribution of 200 posterior draws from the model.

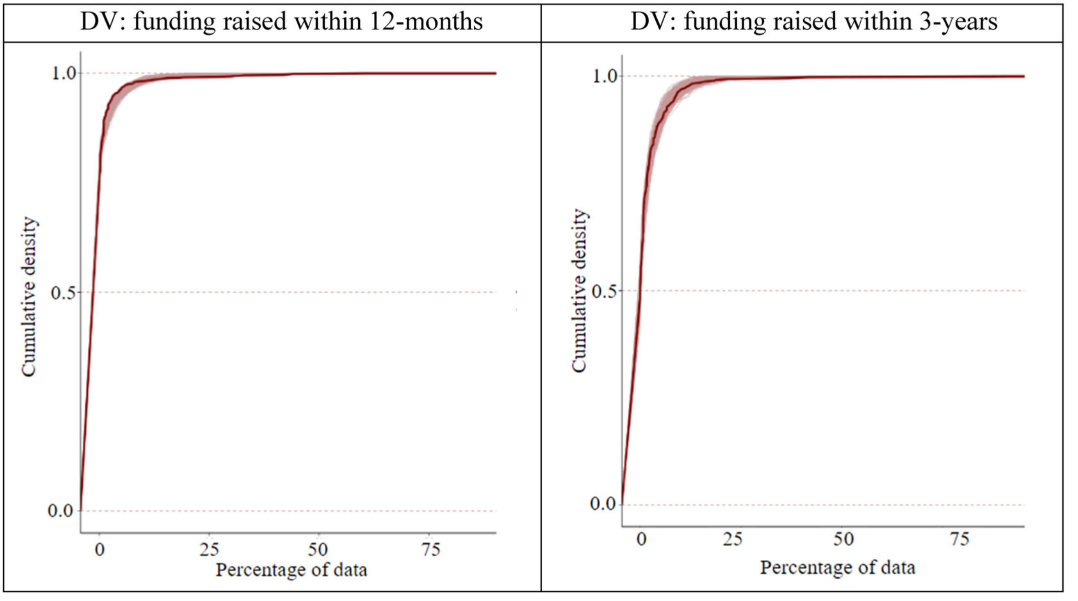

FIGURE B8 Empirical cumulative distribution function (ECDF) overlay of DVs compared to posterior draws. Dark red line: Cumulative density of the dependent variable. Red shaded area: Cumulative density of 200 posterior draws from the model.

dependent variable. The light-shaded areas represent the distribution of 200 posterior draws from the model.[18] The posterior draws of the model closely approximate the DV.

Empirical cumulative distribution function

To further validate model fit, we follow Vehtari et al. (2015) and report the empirical cumulative distribution function (ECDF). Figure B8 shows the ECDF of the dependent variables and 200 draws of the posterior of the model.[19] It shows the *cumulative distribution function* of the posterior draws (red shaded area) in comparison to the DVs (solid red line), thus illustrating whether there are any systematic deviations between the two. As indicated by our density overlay plot, the distributions generated by the models fit well with the distribution of the DVs.

Figure B9 is an extension of the comparison between DVs and distributions generated by the models. It shows an empirical cumulative distribution function–probability integral transformation (ECDF-PIT) (Gelman et al., 2013, chap 6; Gabry et al., 2019). The probability integral transformation transforms the probabilistic forecast of the model into a uniform distribution. If the model is perfectly calibrated the plot should show a 45-degree line. The additional light gray areas indicate whether the model significantly deviates from a 45% line within a 95% confidence interval. Both, the model for the 12-month DV and the model for the 3-year DV do not cross the confidence intervals, therefore showing they do not significantly deviate from a straight line and consequently fit the data well.

[18]It is custom to take less than the full set of posterior draws here, for visualization reasons.

[19]It is custom to take less than the full set of posterior draws here, for visualization reasons.

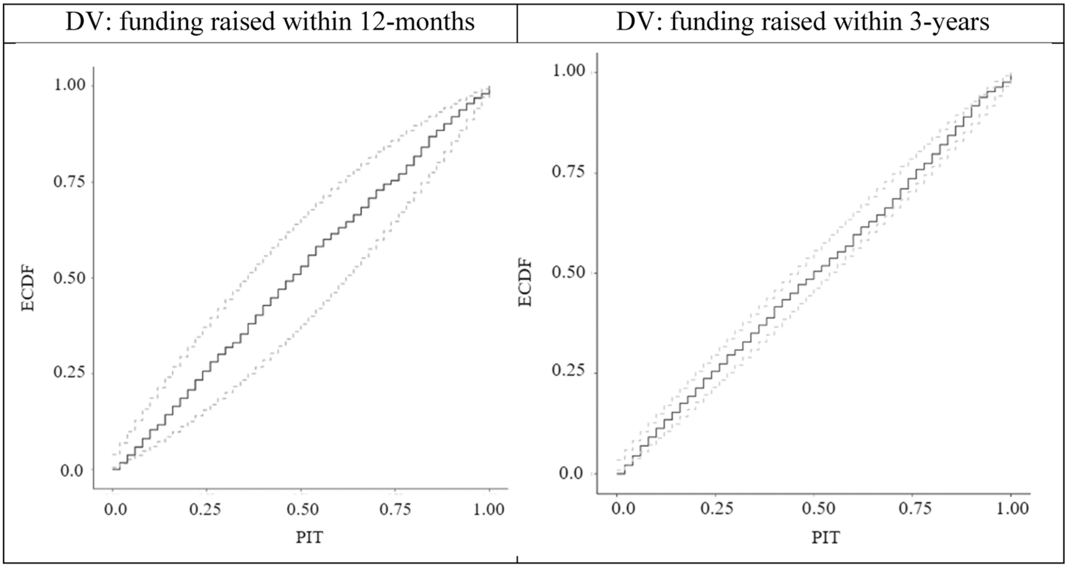

FIGURE B9 Empirical cumulative distribution function–for probability integral transform (ECDF-PIT) plots. Solid black line: Probability integral transformation. Black dotted area: 95% confidence interval.

Holdout sampling

Beyond the fit-based validation methods for Bayesian models presented above, additional analyses can be conducted that involve holdout sampling and give some indication of the predictive validity of the model.

The Pareto *K*-diagnostic leave-one-out cross-validation (LOO-CV) based on Pareto smoothed importance sampling (PSIS), is a Bayesian version of leave-one-out cross-validation (Vehtari et al, 2015). It estimates whether particular data points—if left out of the estimation—would have been predicted well by the model. It consequently identifies data points that the model does not predict well. Thus, while many of the tests described so far assess to which extent the model is an overall good fit for the data, this test assesses the degree to which the model predicts the value of individual data points well.

Tables B2 and B3 below show the distribution of values for Pareto *K*-diagnostics. Most data points have values that are smaller than 0.5, indicating that the leave-one-out cross-validation model predicts these data points well (Vehtari et al, 2015). For the 12-month DV, out of the 1350 data points, there are 1322 (97.9%) data points that the model predicts well (Pareto *K* value (−inf, 0.5]), 19 data points (1.4%) that the model predicts okay (Pareto *K* value (0.5, 0.7]), 9 (0.7%) that it does not predict well (Pareto *K* value (0.7, 1]), and there are no observations that are not predict well (Pareto *K* value (1, Inf)) (Vehtari et al, 2015). For the 3-year DV, out of the 515 data points, the model predicts 489 (95.0%) of the values well (Pareto *K* value (−inf, 0.5]) and the remaining 26 (5.0%) okay (Pareto *K* value (0.5, 0.7]).

PPC-LOO-CV-PIT plots

The posterior predictive checks (PPC) leave-one-out (LOO) cross-validation (CV) probability integral transformation (PIT) shows how well the model predicts by comparing the distribution

TABLE B2 Pareto *K*-diagnostic DV = funding raised within 12-months.

Pareto *K* value	Count	Percentage
(−inf, 0.5]	1322	97.9%
(0.5, 0.7]	19	1.4%
(0.7, 1]	9	0.7%
(1, Inf)	0	0.0%

TABLE B3 Pareto *K*-diagnostic DV = funding raised within 3-years.

	Count	Percentage
(−inf, 0.5]	489	95.0%
(0.5, 0.7]	26	5.0%
(0.7, 1]	0	0.0%
(1, Inf)	0	0.0%

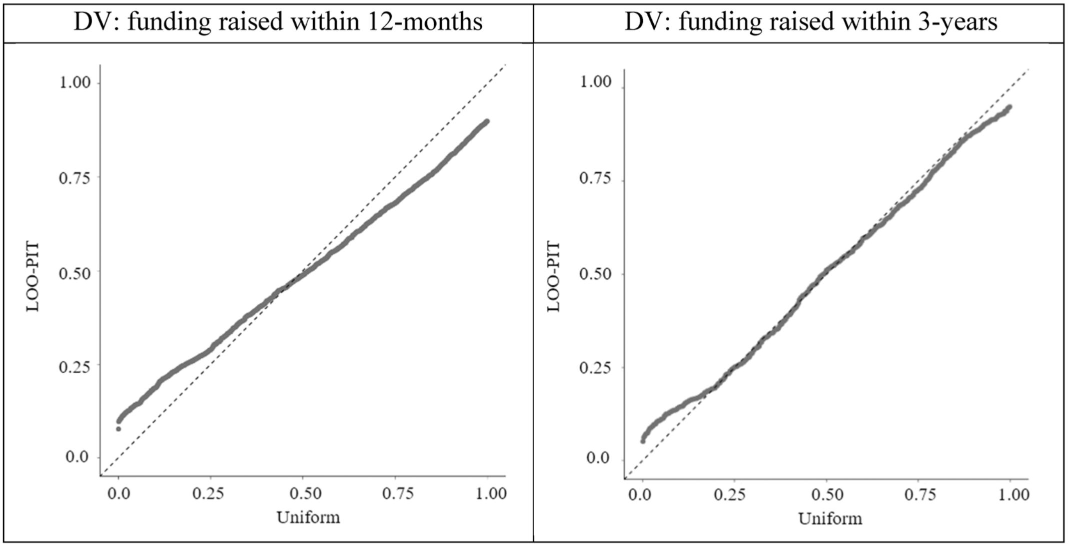

FIGURE B10 Posterior predictive checks (PPC) leave-one-out (LOO) cross-validation (CV) probability integral transformation (PIT) (PPC-LOO-CV-PIT).

of LOO-PIT-CV values to the expected uniform distribution if the model is well-calibrated (Gelman et al., 2013, chap 6; Gabry et al., 2019). A 45-degree line indicates that the model perfectly predicts all data points, while moving from the 45-degree line indicates some deviation from this perfect prediction. On this predictive check, the 12-month model shows a slight deviation, while the 3-year model has an almost perfect predictive performance with minimal deviation (Figure B10).

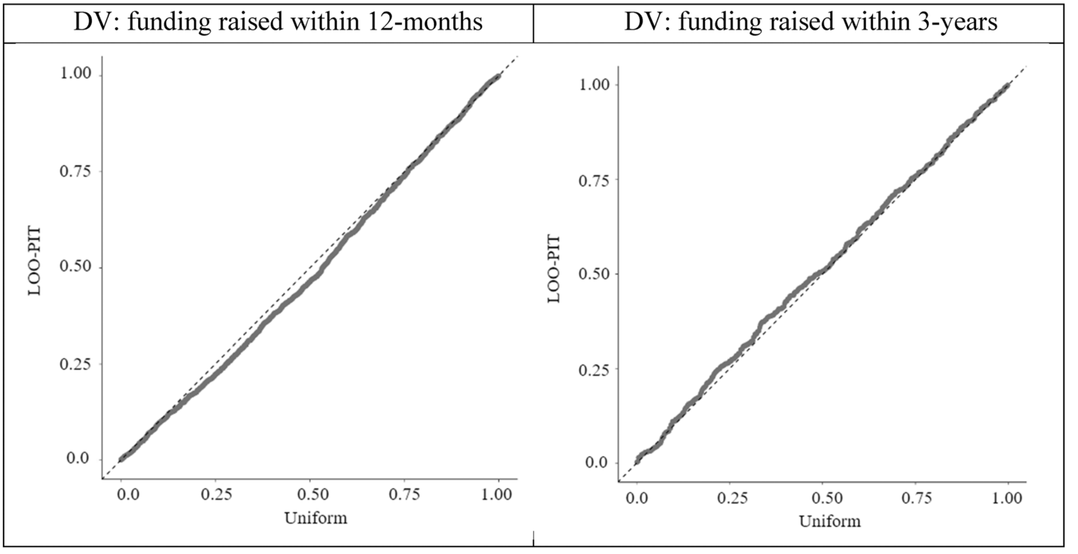

FIGURE B11 Outlier robust PPC-LOO-CV-PIT.

This slight deviation could be driven by outliers. To test this, we run an outlier robust model as developed by Mangiola et al. (2021) (also see the section on additional analyses in the paper).

For the 12-month DV, this approach defines 32 observations as outliers; we then run the model without these outliers to see if doing so affects our effect sizes. The PPC-LOO-CV-PIT for this outlier robust model are presented in Figure B11. The effect sizes can be found in Table C2 below. Figure B11 shows an almost perfect fit. The effect sizes of this model only deviate slightly from our main results and are within the posterior intervals of our main results.

APPENDIX C: ROBUSTNESS TESTS AND ANALYSIS

We conduct additional analyses to check the robustness of the findings across different samples and methodologies. Table C1 provides an overview of the different analyses and the underlying rationale, and Table C2 presents the results of these analyses.

First, as explained above, the Bayesian analysis should be robust to small group sizes; nevertheless, we also conducted an analysis in which our group sizes (number of cells in each low-level group) were larger than in the main analysis—thus avoiding data spareness concerns. To that end we estimate our Bayesian models without the manager effect. Not including the manager in the hierarchical structure leads to larger group sizes: An average of 5.2 and 6.6 cohorts per accelerator for the 12-months and 3-year dependent variables, respectively. Table C2 reports the results. We observe an accelerator effect of 1.5% and a cohort effect of 8.6% for fundraising during the 12 months following the startups' accelerator entry. Similarly, we observe a 1.7%

TABLE C1 Overview of reported analyses.

Analysis	Reason for analysis	Notes
Bayesian hierarchical model with a Tweedie distribution	The method handles *zero-inflated* dependent variables as well as *sparse small "group sizes."*	This is the main methodology. All additional analyses are estimated using it.
Run model without estimating a manger effect	Doing so alleviated data sparseness because it results in larger *group sizes* (cohorts per accelerator).	No substantial change to results.[a]
Delete cohorts where no startup secured funding	Test if the cohort effect is driven by a few startups that secured funding.	No substantial change to results.[a]
Delete cohorts where ≤1 startups secured funding		
Delete cohorts where ≤2 startups secured funding		
Delete cohorts where ≤3 startups secured funding		
Bayesian outlier analyses based on Mangiola et al. (2021)	Test robustness to outliers, defined as data points, not well predicted by the model.	No substantial changes to the results.[a]
Winsorized at mean +3 standard deviations	Test if the cohort effect is driven by one (or more) "star" startups that fundraised a substantial amount of investment.	No substantial changes to the results.[a]
Winsorized at mean +4 standard deviations		
Winsorized above "jump" in DV		
Include number of cohorts as a control	Test if cohort effect is due to accelerator learning to improve from one cohort to another.	No substantial change to results.[a]
Include controls for startup features	Test if effects are driven by differences (between cohorts/accelerators) in startup features (age, number of founders).	No substantial change to results.[a]

[a]We observe small changes to a few effect sizes, yet the results show the same pattern of relative magnitudes as our primary results.

TABLE C2 Comparison of main results to results of additional robustness analyses.

Panel A: The dependent variable is capital raised 12 months after startups' accelerator entry				
	Results main analysis	**No manager effect**	**Delete cohorts where no one fundraised**	**Delete cohorts where ≤1 startup fundraised**
Accelerator	3.60	1.45	1.94	3.96
Manager	7.70	–	8.39	6.73
Cohort	7.51	8.63	8.12	8.66
Industry	3.18	3.23	2.14	1.32
Year	0.00	0.00	0.00	0.01
Pre. funding	0.07	0.08	0.04	0.08
n	1350	1350	1134	939
	Delete cohorts where ≤2 startup fundraised	**Delete cohorts where ≤3 startup fundraised**	**Delete outliers as defined by** Mangiola et al. (2021)	**Winsorized at mean +3 StD**
Accelerator	4.44	2.33	2.66	3.56
Manager	7.59	7.72	7.62	7.10
Cohort	8.86	9.16	8.49	7.64
Industry	1.72	2.72	2.27	3.67
Year	0.01	0.00	0.00	0.00
Pre. funding	0.00	0.05	0.06	0.08
n	736	582	1318	1350
	Winsorized at mean +4 StD	**Winsorized above "natural jump" (23 M)**	**Winsorize at: 20 M**	**Winsorize at: 18 M**
Accelerator	2.68	4.08	3.37	2.80
Manager	8.73	8.01	7.77	7.43
Cohort	8.01	7.77	7.72	7.82
Industry	4.43	4.18	4.38	3.51
Year	0.00	0.00	0.00	0.00
Pre. funding	0.07	0.07	0.08	0.07
n	1350	1350	1350	1350
	Winsorize at: 16 M	**Winsorize at: 14 M**	**Number of cohorts as a control**	**With startup controls**
Accelerator	3.36	3.79	3.62	4.35
Manager	8.07	8.09	7.85	8.47
Cohort	7.05	7.46	7.53	8.15
Industry	4.30	3.78	0.68	2.50
Year	0.00	0.00	0.01	0.00

TABLE C2 (Continued)

Panel A: The dependent variable is capital raised 12 months after startups' accelerator entry				
	Results main analysis	**No manager effect**	**Delete cohorts where no one fundraised**	**Delete cohorts where ≤1 startup fundraised**
Pre. funding	0.06	0.06	0.07	0.08
n	1350	1350	1350	1350

Panel B: The dependent variable is capital raised 3 years after accelerator entry				
	Results main analysis	**No manager effect**	**Delete cohorts where no startup fundraised**	**Delete cohorts where ≤1 startup fundraised**
Accelerator	2.60	1.72	2.49	2.23
Manager	4.91	–	5.18	4.99
Cohort	5.64	8.01	5.13	5.84
Industry	1.56	0.74	1.28	1.46
Year	0.00	0.00	0.00	0.00
Pre. funding	0.00	0.01	0.01	0.00
n	515	515	488	427
	Delete cohorts where ≤2 startup fundraised	**Delete cohorts where ≤3 startup fundraised**	**Delete outliers as defined by** Mangiola et al. (2021)	**Winsorize at mean +3 StD**
Accelerator	2.70	3.05	2.35	2.29
Manager	5.28	5.21	4.95	4.88
Cohort	5.58	5.81	5.17	4.94
Industry	1.12	1.32	0.86	1.47
Year	0.00	0.00	0.00	0.00
Pre. funding	0.00	0.00	0.00	0.00
n	378	339	491	515
	Winsorize at mean +4 StD	**Winsorize above "natural jump" (26 M)**	**Winsorize at: 20 M**	**Winsorize at: 18 M**
Accelerator	2.14	2.46	2.34	2.79
Manager	5.51	4.93	5.08	5.62
Cohort	5.13	5.51	5.14	5.43
Industry	2.22	1.18	1.42	1.72
Year	0.00	0.00	0.00	0.00
Pre. funding	0.00	0.00	0.00	0.00
n	515	515	515	515

TABLE C2 (Continued)

Panel B: The dependent variable is capital raised 3 years after accelerator entry				
	Results main analysis	No manager effect	Delete cohorts where no startup fundraised	Delete cohorts where ≤1 startup fundraised
	Winsorize at: 16 M	Winsorize at: 14 M	Number of cohorts as a control	With startup controls
Accelerator	1.89	2.61	1.87	1.77
Manager	5.33	5.16	5.13	6.21
Cohort	5.52	4.94	5.43	5.80
Industry	1.39	1.27	0.49	1.36
Year	0.00	0.00	0.00	0.00
Pre. funding	0.00	0.00	0.00	0.00
n	515	515	515	515

accelerator effect and an 8.0% cohort effect for the 3-year fundraising dependent variable. These findings document a sizable cohort effect and suggest that the cohort effect reported in the primary analysis is unlikely to be an artifact of data sparseness.

Next, we examined whether our results—especially the cohort effect introduced in this study—are driven by cohorts that have only one or a handful of startups that raised funds.[20] That is, we ask if the cohort effect is driven by entire cohorts or if it simply captures one or a few startups within a cohort that successfully fundraised. To address this concern, we conducted several analyses where we removed all cohorts in which no startups, 1 or fewer, 2 or fewer, and 3 or fewer, received an investment. Doing so results in similar effect sizes as our main analysis (see Table C2). We believe these analyses show that the effects are not merely driven by a handful of "star" startups and thus underscore our call to better understand bottom-up cohort dynamics.

We also test for the potential impact of outlier values of the two DVs in multiple ways. First, we follow Mangiola et al. (2021), who define outliers not based on some value threshold but based on a Bayesian model's predictive abilities. In their approach "new data are generated from the fitted model, providing the theoretical range of values for each data point. All observed read counts that are outside the 95% posterior credible interval are quarantined as possible outliers" (Mangiola et al., 2021, p. 3). For the 12-month DV, this approach defines 32 observations, and for the 3-year DV, 24 observations as outliers. We then run the model without these outliers to see if doing so affects our effect sizes. For both DVs effect sizes of this model are within the posterior intervals of our main results.

We also conducted a sensitivity analysis where we winsorize our dependent variables at 3 and 4 standard deviations above the mean and we visually inspected the data for any "natural jumps" in the DV above which there are only a handful of outliers and winsorized the data above the jump.[21] For example, we observe such a jump at $23 million for the 12-month

[20]We thank an anonymous reviewer for this suggestion.

[21]We thank the anonymous reviewers for this suggestion.

fundraising amount ($26 million for the 3-year fundraising amount). To examine if our main results are driven by a few startups that gain large funding amounts even further, we then reduced the winsorizing points (in both samples) to $20, $18, $16, and $14 million. Doing so reduces the magnitude of the cohort effect marginally, which is to be expected in the absence of large values for the DV. But across all models (in both samples), the magnitude of the cohort effect remains larger than the accelerator effect. These additional analyses confirm that the notable cohort and manager effects are not driven by outliers.

The battery of analyses reported above addresses questions about the method as well as the features of the data. Below, we report two additional analyses that are more conceptual in nature. First, accelerators may learn how to improve their top-down activities from cohort to cohort. As a result, the cohort effect we attributed to bottom-up mechanisms may simply reflect a within-accelerator improvement over time. We test for this explanation by including the number of previous cohorts before a focal one as a control in our model; doing so does not result in a substantial change to the cohort effect. The analysis suggests that the cohort effect is not merely an artifact of accelerators' improvement via executing top-down mechanisms.

Second, we recognize that some of the variance in startup performance is due to innate startup traits. It could be informative to estimate the degree to which differences in the startups themselves explain differences in funding outcomes. Unfortunately, it is not possible to include a startup effect in our variance-decomposition analysis. This is because the data is at the startup level (i.e., it includes a single observation for each startup), while estimating a variance-decomposition effect requires the effect to be associated with more than one observation (otherwise the model would be over-specified). In the absence of the ability to include a startup effect, we run the analysis using observable features of the startups: *Age* at time of entry into the accelerator and *number of founders*. Doing so does not lead to a substantial change in the magnitudes of the accelerator, manager and cohort effects.

APPENDIX D: CODE FOR THE BAYESIAN HIERARCHICAL ANALYSIS (STAN & R)

CMDSTANR model

```
functions {
  int num_non_zero_fun(vector y) {
    int A = 0;
    int N = num_elements(y);
    for (n in 1 : N) {
      if (y[n] != 0) {
        A += 1;
      }
    }
    return A;
  }

  array[] int non_zero_index_fun(vector y, int A) {
    int N = num_elements(y);
    array[A] int non_zero_index;
    int counter = 0;
    for (n in 1 : N) {
      if (y[n] != 0) {
        counter += 1;
        non_zero_index[counter] = n;
      }
    }
    return non_zero_index;
  }

  array[] int zero_index_fun(vector y, int Z) {
    int N = num_elements(y);
    array[Z] int zero_index;
    int counter = 0;
    for (n in 1 : N) {
      if (y[n] == 0) {
        counter += 1;
        zero_index[counter] = n;
      }
    }
    return zero_index;
  }

  void check_tweedie(real mu, real phi, real mtheta) {
    if (mu < 0) {
      reject("mu must be >= 0; found mu =", mu);
    }
    if (phi < 0) {
```

```
    reject("phi must be >= 0; found phi =", phi);
  }
  if (mtheta < 1 || mtheta > 2) {
    reject("mtheta must be in [1, 2]; found mtheta =", mtheta);
  }
}

void check_tweedie(vector mu, real phi, real mtheta) {
  int N = num_elements(mu);
  if (phi < 0) {
    reject("phi must be >= 0; found phi =", phi);
  }
  if (mtheta < 1 || mtheta > 2) {
    reject("mtheta must be in [1, 2]; found mtheta =", mtheta);
  }
  for (n in 1 : N) {
    if (mu[n] < 0) {
      reject("mu must be >= 0; found mu =", mu[n], "on element", n);
    }
  }
}

real tweedie_lpdf(vector y, vector mu, real phi, real mtheta, int M) {
  check_tweedie(mu, phi, mtheta);
  int N = num_elements(y);
  int N_non_zero = num_non_zero_fun(y);
  int N_zero = N - N_non_zero;
  array[N_zero] int zero_index = zero_index_fun(y, N_zero);
  array[N_non_zero] int non_zero_index = non_zero_index_fun(y, N_non_zero);
  int A = num_elements(non_zero_index);
  int NmA = N - A;
  vector[N] lambda = 1 / phi * mu ^ (2 - mtheta) / (2 - mtheta);
  real alpha = (2 - mtheta) / (mtheta - 1);
  vector[N] beta = 1 / phi * mu ^ (1 - mtheta) / (mtheta - 1);
  real lp = -sum(lambda[zero_index]);
  for (n in 1 : A) {
    vector[M] ps;
    for (m in 1 : M) {
      ps[m] = poisson_lpmf(m | lambda[n])
              + gamma_lpdf(y[non_zero_index[n]] | m * alpha, beta[n]);
    }
    lp += log_sum_exp(ps);
  }
  return lp;
}

real tweedie_rng(real mu, real phi, real mtheta) {
  check_tweedie(mu, phi, mtheta);
  real lambda = 1 / phi * mu ^ (2 - mtheta) / (2 - mtheta);
```

```
    real alpha = (2 - mtheta) / (mtheta - 1);
    real beta = 1 / phi * mu ^ (1 - mtheta) / (mtheta - 1);
    int N = poisson_rng(lambda);
    real tweedie_val;
    if (mtheta == 1) {
      return phi * poisson_rng(mu / phi);
    }
    if (mtheta == 2) {
      return gamma_rng(1 / phi, beta);
    }
    if (N * alpha == 0) {
      return 0.;
    }
    return gamma_rng(N * alpha, beta);
  }
}

data {
  int<lower=1> N; // total number of observations
  vector[N] Y; // response variable
  int<lower=1> K; // number of population-level effects
  matrix[N, K] X; // population-level design matrix
  // data for group-level effects of ID 1
  int<lower=1> N_1; // number of grouping levels
  int<lower=1> M_1; // number of coefficients per level
  array[N] int<lower=1> J_1; // grouping indicator per observation
  // group-level predictor values
  vector[N] Z_1_1;
  // data for group-level effects of ID 2
  int<lower=1> N_2; // number of grouping levels
  int<lower=1> M_2; // number of coefficients per level
  array[N] int<lower=1> J_2; // grouping indicator per observation
  // group-level predictor values
  vector[N] Z_2_1;
  // data for group-level effects of ID 3
  int<lower=1> N_3; // number of grouping levels
  int<lower=1> M_3; // number of coefficients per level
  array[N] int<lower=1> J_3; // grouping indicator per observation
  // group-level predictor values
  vector[N] Z_3_1;
  // data for group-level effects of ID 4
  int<lower=1> N_4; // number of grouping levels
  int<lower=1> M_4; // number of coefficients per level
  array[N] int<lower=1> J_4; // grouping indicator per observation
  // group-level predictor values
  vector[N] Z_4_1;
  int prior_only;
  int<lower=1> M;
}
```

```
transformed data {
}

parameters {
  vector[K] b; // population-level effects
  real<lower=0> phi;
  real<lower=1, upper=2> mtheta;

  vector<lower=0>[M_1] sd_1; // group-level standard deviations
  array[M_1] vector[N_1] z_1; // standardized group-level effects

  vector<lower=0>[N_1] sd_2; // group-level standard deviations - N_1 is the number of groupings of accelerators
  array[M_2] vector[N_2] z_2; // standardized group-level effects

  vector<lower=0>[N_2] sd_3; // group-level standard deviations -  N_2 is the number of groupings of managers
  array[M_3] vector[N_3] z_3; // standardized group-level effects

  vector<lower=0>[M_4] sd_4; // group-level standard deviations
  array[M_4] vector[N_4] z_4; // standardized group-level effects
}

transformed parameters {
  vector[N_1] r_1_1; // actual group-level effects
  vector[N_2] r_2_1; // actual group-level effects
  vector[N_3] r_3_1; // actual group-level effects
  vector[N_4] r_4_1; // actual group-level effects
  real lprior = 0;
  r_1_1 = sd_1[1] * z_1[1];
  r_2_1 = sd_2[1] * z_2[1];
  r_3_1 = sd_3[1] * z_3[1];
  r_4_1 = sd_4[1] * z_4[1];
  lprior += gamma_lpdf(phi | 0.01, 0.01);
  lprior += student_t_lpdf(sd_1 | 3, 0, 2.5)[1]
          - 1 * student_t_lccdf(0 | 3, 0, 2.5);
  lprior += student_t_lpdf(sd_2 | 3, 0, 2.5)
          - 1 * student_t_lccdf(0 | 3, 0, 2.5);
  lprior += student_t_lpdf(sd_3 | 3, 0, 2.5)
          - 1 * student_t_lccdf(0 | 3, 0, 2.5);
  lprior += student_t_lpdf(sd_4 | 3, 0, 2.5)
          - 1 * student_t_lccdf(0 | 3, 0, 2.5);
}

model {
  // likelihood including constants
  if (!prior_only) {
    // initialize linear predictor term
```

```
    vector[N] mu = rep_vector(0.0, N);
    mu += Intercept + X * b;
    for (n in 1 : N) {
      // add more terms to the linear predictor
      mu[n] += r_1_1[J_1[n]] * Z_1_1[n] + r_2_1[J_2[n]] * Z_2_1[n]
               + r_3_1[J_3[n]] * Z_3_1[n] + r_4_1[J_4[n]] * Z_4_1[n];
    }
    target += tweedie_lpdf(Y | mu, phi, mtheta, M);
  }
  // priors including constants
  target += lprior;
  target += std_normal_lpdf(z_1[1]);
  target += std_normal_lpdf(z_2[1]);
  target += std_normal_lpdf(z_3[1]);
  target += std_normal_lpdf(z_4[1]);
}

generated quantities {
  // actual population-level intercept
  real b_Intercept = Intercept;
  vector[N] mu = rep_vector(0.0, N);
  mu += Intercept + X * b; // Adding the intercept and fixed effects
  for (n in 1 : N) {
    // Adding group-level effects
    mu[n] += r_1_1[J_1[n]] * Z_1_1[n] + r_2_1[J_2[n]] * Z_2_1[n]
             + r_3_1[J_3[n]] * Z_3_1[n] + r_4_1[J_4[n]] * Z_4_1[n];
  }

  // Generating simulated data based on the model
  vector[N] r_tweedie;
  for (n in 1:N) {
    r_tweedie[n] = tweedie_rng(mu[n], phi, mtheta);
  }
}

fit_stan <- model$sample(
  data = stan_data,
  chains = 2,           # Number of Markov chains
  parallel_chains = 2,  # Number of chains to run in parallel
  iter_warmup = 2000,   # Number of warmup iterations per chain
  iter_sampling = 2000, # Number of sampling iterations per chain
)
```