NBER WORKING PAPER SERIES

CHOOSING TECHNOLOGY:
AN ENTREPRENEURIAL STRATEGY APPROACH

Joshua S. Gans
Michael Kearney
Erin L. Scott
Scott Stern

Working Paper 27489
http://www.nber.org/papers/w27489

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2020

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w27489.ack

Choosing Technology: An Entrepreneurial Strategy Approach
Joshua S. Gans, Michael Kearney, Erin L. Scott, and Scott Stern
NBER Working Paper No. 27489
July 2020
JEL No. O31

ABSTRACT

A central premise of research in the strategic management of innovation is that start-ups are able to leverage emerging technological trajectories as a source of competitive advantage. But, if the potential for a technology is given by the fundamental character of a given technological trajectory, then why does entrepreneurial strategy matter? Or, put another way, if the evolution of technology is largely shaped by the strategic choices entrepreneurs make, then why do technological trajectories exhibit systematic patterns such as the Technology S-curve? Taking a choice-based perspective, this paper illuminates the choices confronting a start-up choosing their technology by resolving the paradox of the Technology S-curve through a reformulation of the foundations of the Technology S-curve. Specifically, we reconceptualize the Technology S-curve not as a technological given but as an envelope of potential outcomes reflecting differing strategic choices by the entrepreneur in exploration versus exploitation. Taking this lens, we are able to clarify the role of technological uncertainty on start-up strategy, the impact of constraints on technological evolution, and how technology choice is shaped by the possibility of imitation. Our findings suggest that staged exploration may stall innovation as a result of the replacement effect, increasing the strategic importance of commitment.

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
CANADA
and NBER
joshua.gans@gmail.com

Michael Kearney
Massachusetts Institute of Technology
Sloan School of Management
100 Main Street, E62-463
Cambridge, MA 02142
mkearney@mit.edu

Erin L. Scott
MIT Sloan School of Management
100 Main St
E62-470
Cambridge, MA 02142
United States
elscott@mit.edu

Scott Stern
MIT Sloan School of Management
100 Main Street, E62-476
Cambridge, MA 02142
and NBER
sstern@mit.edu

I. Introduction

Technological trajectories are structured yet evolutionary processes by which emerging technologies change and improve over time (Dosi 1982), with implications for the distribution of value from innovation (Teece 1986), the co-evolution of industries and competition (Utterback 1994), and the choice of commercialization strategy (Teece 1986; Gans and Stern 2003). The analytical insights that arise from these studies are grounded in the thesis that, even in the face of significant uncertainty, the evolution of technology follows a structured process, and that managers (and, especially entrepreneurs) may be able to harness this process as a foundation for competitive advantage. This perspective is perhaps most forcefully articulated through the now-pervasive concept of the Technology S-curve (Foster 1986). Foster argued that, at the earliest stage of an emerging technological trajectory, investment in innovation yields few concrete performance improvements; realized research productivity is dampened by learning, experimentation, and failures. But, for those technologies that are ultimately impactful, an inflection point is reached, at which firms focus on exploitation of their learning, R&D productivity increases dramatically, and only tapers off as the technology reaches certain physical or architectural limits. This concept of S-shaped technological change is central across a wide swath of research in the management of innovation, including the pioneering work of Henderson and Clark (1990) on the role of architectural innovation, as well as the influential work by Christensen (1997) and his theory of disruptive innovation.

Though these foundational studies focused on the role of and implications for incumbent firms, they also offered an important lens for entrepreneurs evaluating potential opportunities from an emerging technology. Many an entrepreneur (knowingly or not) has built their entrepreneurial strategy on theories, empirics, and stylized facts grounded in our understanding of the Technology S-curve. While the ultimate impact of a technological trajectory depends on more than its ability to achieve a given level of technological performance (Adner and Kapoor 2016), it is nonetheless true that the ability of entrepreneurs to achieve meaningful technological advance with an emergent technology offers the potential for that firm to achieve competitive advantage. Free from the organizational inertia that may bias existing firms and delay their engagement with an emerging trajectory (Ganco and Agarwal 2009; Wu et al 2013; Kapoor and Furr 2015), entrepreneurs must choose and quickly develop robust organizational capabilities to realize the competitive advantage this freedom offers (Mitchell 1994; Carroll et al 1996; Klepper and Simons 2000; Khessina and Carroll 2008). These early strategic choices, including technological investments, are critical to not only realizing their competitive advantage, but also their long-term survival (Furr et al 2012).

At the heart of the Technology S-curve lies the fact that successful innovation requires both investments in exploration and exploitation, a central insight within innovation strategy (and organizational learning more generally (March 1991)). Exploiting existing technology in the near term offers an immediate advantage but may imperil long-term survival and growth, while exploring novel pathways opens up routes to long-term success but may involve ceding near-term opportunities (Levinthal and March 1981; March 1991). This tension is front of mind for entrepreneurs. Near term opportunities to exploit an emerging technology may facilitate key learnings (for instance, from customers) and the acquisition of resources necessary to enable a longer period of exploration; yet, these exploitation activities may distract from (and potentially constrain, where commitments are involved) the long-term objectives of the start-up. Indeed, the empirical instantiation of the Technology S-curve as a phenomenon is perhaps the most commonly invoked exemplar of the tradeoff between exploration and exploitation.

But, these two arguments rest on a seeming paradox. Specifically, while on the one hand, the Technology S-curve is taken as an empirical phenomenon that at its essence is technologically determined, the argument justifying the shape of the S-curve depends on strategic choices as to the balance between exploration and exploitation. But, if S-curves are an environmental fact, then why do these strategic choices matter? Alternatively, if they are the result of strategic choice, why do they appear as a systematic empirical phenomenon? Said otherwise, is the role of entrepreneurial strategy to inform entrepreneurs as they select and then 'ride' an emerging technology trajectory? Or, does a robust entrepreneurial strategy necessitate understanding how the strategic choices of the entrepreneur will shape the emerging technology trajectory itself?

Consider the case of the electric vehicle industry in the late 2000s, and the contrasting strategies of start-ups Better Place and Tesla. On the one hand, both Better Place and Tesla premised their innovation strategy on the potential application of lithium ion batteries, which were widely perceived to be undergoing significant technological advancement. Within the domain of lithium ion batteries, though, Better Place founder Shai Agassi focused quickly on the potential exploitation of existing lithium ion battery technology. While this still emerging technology had significant limitations including a high cost and low energy storage (and so vehicle range would be limited and charging times were long), Agassi prioritized the development of a "swappable" battery in which customers could simply replace the batteries in a few minutes at a "replacement" station rather than necessarily having to wait while the battery itself was being charged. This pragmatic near-term focus allowed Better Place to establish partnerships with traditional automakers, finance regional infrastructure (e.g., swapping stations across the entirety of Israel), and introduce a "mainstream" electric vehicle by 2012 (Squatriglia 2009).

In contrast, Elon Musk recognized the same broad opportunity and selected the same emerging technological trajectory (lithium ion batteries). However, in contrast to Better Place, Tesla pursued a fundamentally different innovation management approach, prioritizing exploration over exploitation, with two key features. First, Musk implemented a deliberative stepwise model of introduction. Accounting for the existing shortcomings of the still nascent battery technology, Tesla introduced the Roadster, a luxury vehicle, where price was not a key factor in the purchase decision. The experience developing the Roadster enabled Tesla to explore continued development of the lithium ion battery technology, take advantage of learning-by-doing in process development, and ultimately benefit from economies of scale (e.g., through vertical integration into battery manufacturing at the Giga Factory) (Baker 2016). From 2010 through 2018, lithium-ion battery costs declined 90 percent, enabling increases in energy, power density and vehicle range.

Though both Better Place and Tesla pursued the "same" emerging technological opportunity (at the same time, industry, and domain of application), their ability to leverage the emerging technological innovation as a source of competitive advantage diverged. After an initial rollout in 2012 that met with a very low level of demand (fewer than 1500 Better Place vehicles were sold), Better Place was unable to shift away from its exploitation-oriented approach to take advantage of the emerging improvements being developed by Tesla (Chafkin 2014). In contrast, Tesla's long-term strategy of pairing a more exploratory approach with a focus on the luxury vehicle segment allowed it to develop and commercialize the "mainstream" Model 3 sedan by 2018. Despite more than $1 billion USD in investment in Better Place, the entire venture was liquidated for scrap value in 2013, while Tesla was able to "ride the S-curve" to emerge as the leading electric car manufacturer for advanced economies such as the United States and Europe.

The contrasting experiences of Tesla and Better Place puts into sharp relief both the central insight but also the implicit paradox behind the Technology S-curve. On the one hand, the relative success of Tesla seems to be a classic case where an innovation strategy premised on pioneering a new S-curve allows a start-up venture to ultimately compete with more established players by taking advantage of improved performance in a nascent technological trajectory over time. At the same time, despite investing in the same technology (and market, etc.), the choice of Better Place to "exploit" the new S-curve more quickly seems to have precluded its ability to take advantage of the more exploratory path pioneered by Tesla. What is the relationship between the Technology S-curve traced out by Better Place (low exploration and high exploitation) versus Tesla (with a longer and more intensive period in the exploration phase)? And, how does the fact that the S-curve that each start-up followed reflected the strategic choices by each founding team impact how we interpret and identify Technology S-curves themselves?

The purpose of this paper is to better illuminate the choices confronting a start-up choosing their technology by resolving the paradox of the Technology S-curve through the reformulation of the foundations of the Technology S-curve. While most research in innovation and strategy treats the Technology S-curve as a given (e.g., a technological trajectory is likely to be disruptive or not), we reformulate technology strategy through a choice-based framework in which the S-curve realized by a given start-up depends on the choices of its founders. We reconceptualize the technology S-curve, not as a technological given, but as an envelope of potential outcomes derived by strategic choice, and we explicitly account for the active choice technology entrepreneurs make. This perspective suggests that the limiting factor shaping long-term innovation performance is not fundamental physical limits but strategy itself.

This choice-based approach offers novel insight into the strategic implications of the tradeoff between exploration and exploitation. Whereas the explore-exploit framework has traditionally reflected strategic choices by firms choosing between alternative and distinct Technology S-curves (Foster 1986; Anderson and Tushman 1990; Henderson and Clark 1990; Christensen 1997), our framework offers a complementary strategic lens (most relevant for resource-constrained entrepreneurs) about how to engage with a single but still emerging technological trajectory. Rather than focusing on how to divide effort and attention *across* technologies (e.g., the choice by General Motors to prioritize internal combustion versus battery-powered vehicles), our framework emphasizes the importance of active choice in shaping the evolution of the emerging technology itself (e.g., the choice of *when* to take a more exploratory versus more exploitation-oriented focus towards lithium-ion battery technology for automobiles).

This additional choice-based perspective is important insofar as a significant literature assumes (often implicitly) that a certain group of technologies (e.g., lithium ion batteries) have the potential for disruption of the established market; however, a choice-based approach emphasizes that whether a technology ends up being disruptive depends less on the inherent nature of the technology but on whether the entrepreneur chooses a strategy in order to be disruptive to established players. Put another way, observations of disruptive innovation (e.g., overtaking an established trajectory and doing so in a way that overturns established market power) fundamentally conflates the potential for disruption with the strategic choice to take advantage of its disruptive potential; rather than seeking to identify a "disruptive technology," entrepreneurs should consider how their strategic choices are coherent with a disruptive strategy. Indeed, a choice-based approach illuminates the potential for entrepreneurs, particularly those embracing a disruptive strategy, to become trapped by a novel form of the replacement effect whereby success with limited exploration may reduce their incentives to explore the full potential of a novel technological trajectory (and, in fact, disrupt the incumbent).

We begin by considering the nature and history of the Technology S-curve, and the central role that this tool plays (often implicitly) in research, teaching and practice within the broad fields of innovation strategy and technology-driven entrepreneurship. This analysis highlights the "paradox" at the heart of the

Technology S-curve, motivating the choice-based approach for start-ups we develop in Section III. We then present a simple model of start-up technology strategy in Section IV that, though highly stylized, clarifies the role of technological uncertainty on start-up technology strategy, the impact of constraints (such as limited resources) on technological evolution, and how technology choice is shaped by the possibility of imitation. In Section V we consider how the choices of start-up founders to explore or exploit an emerging technology may manifest in the industry-level Technology-curve. While the broad conceptual framing applies to all firms that engage with a new technological trajectory, the strategic tradeoffs we describe are perhaps most salient in the case of a start-up innovator, where both the choice of technological trajectory and choice of exploration (relative to exploitation) may serve as core choices of their entrepreneurial strategy (Gans Scott and Stern 2018; Gans Stern and Wu 2019).

II. The Technology S-Curve Paradox

The Technology S-curve occupies a critical yet elusive position at the intersection between strategy and innovation. At its core, the Technology S-curve offers an elegant and structured characterization of the dynamic evolution of technology. Specifically, the Technology S-curve captures, within a single framework, the role of learning and experimentation as a foundation for subsequent technological improvement, the cumulative nature of technological innovation, and the constraints of design and natural limits on potential technological performance (Foster 1986; Gans 2016). Moreover, the framework has both empirical content (an S-curve can be measured (see, for example, Nagy et al. (2013))) as well as a range of managerial implications (including how to manage along a given S-curve and the opportunities and challenges associated with “switching” S-curves). And, importantly, the empirical existence of Technology S-curves is the foundational building block in key theoretical arguments regarding the organizational and competitive dynamics accompanying technological transitions (Henderson and Clark 1990; Christensen 1993).

But, a fundamental paradox lies at the heart of the Technology S-curve: while the strategic implications of S-curves depend on their empirical existence, the empirical existence of S-curves depends on strategic choice. If S-curves reflect an inevitable empirical pattern, managers can use that knowledge to plan for how to manage innovation at different points along an S-curve and be proactive in switching S-curves in order to maintain technological leadership. At the same time, the rationale for the S-curve itself depends on a strategic choice to focus at the earliest stages on exploration and then turning towards exploitation in order to improve realized performance. But, if their very existence depends on strategic choice, how can we take their existence as given in order to justify strategic choice itself?

Not simply a minor point of potential inconsistency, this seeming contradiction is grounded in the intellectual synthesis that the S-curve represents. On the one hand, building on a long line of research going back (at least) to Schumpeter (1942), researchers in strategy and innovation grappled during the 1970s and 1980s with how to classify technological change in a way that reflected the dynamic and cumulative nature of innovation while also accommodating the presence of significant uncertainty and path dependence (Abernathy and Utterback 1978; Rosenberg 1976; Sahal 1981). Dosi (1982) offered a way through this impasse through the development of the concept of technological trajectories, in which most technological innovation follows a structured pattern of innovating within a particular trajectory, with the understanding that this type of innovation is fundamentally different from the relatively rarer phenomena of innovation that creates a new trajectory. By clarifying the distinction between continuous versus discontinuous innovation, the technological trajectory framework offered guidance for management and policy in which the optimal strategy surrounding a particular innovation depended on whether that innovation represented the evolution of innovation within a trajectory or the establishment of an entirely new trajectory.

While the concept of technological trajectories allowed for the dynamic examination of how innovation evolves within a given technological domain, a parallel literature grappled with how to classify alternative approaches towards the management of the innovation process itself (Burns and Stalker 1961; Nelson 1962; Abernathy 1976). These studies emphasized the wide range of different organizational approaches to innovation (e.g., highly centralized versus widely dispersed authority, high-powered versus low-powered incentives), and the potential importance of aligning innovation strategy and management with the innovation environment (e.g., the use of a fluid approach in an environment with higher uncertainty). Beyond the identification of broad categories and contingencies, a central insight from these studies is the identification of the tradeoff between exploration and exploitation in the innovation process (Levinthal and March 1981; Roberts 1988; March 1991). Whereas exploration is a search process whose aim is to learn about uncertain domains, exploitation is a selection process whose aim is to leverage accumulated knowledge. This tension between exploration and exploitation offers a foundational managerial insight: to the extent that organizational processes (e.g., incentives, routines) reward exploitation (at least in the short term), organizations might underinvest in exploration (even for a given level of appropriability) resulting in long-term loss in organizational learning and performance (Levinthal and March 1993).

The key insight of Foster (1986) was to bring together these two lines of research into a single framework that simultaneously offered structured predictions about the evolution of technology while also offering guidance into technology management. Specifically, at the earliest stages of a new technological trajectory (a la Dosi), the primary focus will be on exploration in order to maximize learning, identify potential dead ends, and gain a more structured understanding of the potential for improvement for a new technology. An important consequence of this learning orientation is that measured technological performance is unlikely to improve significantly during this early period (i.e., a focus on learning is the microfoundation for the prediction that the early stages of a technological trajectory demonstrate a relatively flat performance-effort profile). This accumulated knowledge pays off when the innovator shifts towards exploitation of the knowledge they have gained; from an empirical perspective, this period is marked by an inflection point in which there are sharper gains in measured performance for a given expenditure of R&D effort. But, this focus on exploitation ultimately depletes the technological performance improvement that can be realized from an initial level of technological knowledge. From an empirical perspective, this latter stage is characterized by diminishing returns and a decline in R&D productivity (i.e., the relationship between effort and performance once again flattens out).

Before considering the strategic implications of this framework, it is useful to emphasize that Foster's formulation of the S-curve was not simply conceptual but also empirical in nature. Foster emphasized not simply the broad shape of the S-curve but the potential for concrete empirical analysis on both a retrospective and prospective basis (and was work that was conducted at McKinsey and other consulting firms at the time). In his work, Foster highlighted the empirical nature of the S-curve in a disparate range of settings from the artificial heart to polymer chemicals to tire cords. For example, his observation that the rate of improvement per unit of effort in the DRAM industry prompted Foster to forecast a shift in the locus of competitive advantage away from DRAM manufacturers and towards software manufacturers (a prediction which was ultimately borne out).

Though its origins lay in the domain of technology management and measurement, the primary application of the S-curve has been in the domain of innovation strategy. The central strategic insight implied by the existence of the S-curve is that whereas established firms have comparative advantage in innovation along a technological trajectory, start-ups may be able to gain competitive advantage by pioneering a new Technology S-curve. An immediate consequence of differential investment and commitment to alternative technological trajectories by incumbents versus entrants is the existence of a potential mechanism by which technological innovation serves as a source of creative destruction by overturning existing sources of market power.

Motivated by this insight, a wave of research in the late 1980s and early 1990s sought to characterize how and when innovation gave rise to creative destruction. Most notably, the influential dissertations of Rebecca Henderson (1988) and Clay Christensen (1992a) (both at the Harvard Business School) used detailed industry-level case studies (of the semiconductor photolithographic equipment industry and the disk drive industry, respectively) as a foundation for conceptualizing the conditions under which established firms were able to achieve technological transitions. Henderson and Clark (1990) articulated the ways in which technologies that involved a shift in design (i.e., architecture) that might impact multiple otherwise siloed areas of competency within an established firm pose fundamental organizational challenges for those firms. Specifically, the challenge for incumbents arising from architectural innovation results from the inability to recognize and communicate across the organization about the architectural potential of the emerging technology, and instead focusing on how the innovation reinforces component-level competencies. In turn, Christensen (1993) focused on market-oriented challenges associated with emerging technological trajectories, in which the potential reluctance to adopt innovation due to the potential for cannibalization (Arrow 1962) is reinforced when existing customers of the established firm do not perceive the potential value of the emerging technological trajectory. Specifically, disruptive innovation is characterized by an initial set of performance attributes not valued by the incumbent existing customers (though, importantly valued by other customer groups) and rapid performance improvement along the attributes existing customers valued, ultimately overtaking the performance of the existing technological trajectory. In other words, even when an established firm may be able to accommodate the innovation in an organizational sense, there may be resistance to the innovation based on a low level of initial technological performance.

In both of these influential bodies of work, the emergence of a specific type of novel technological trajectory has the consequence of offering a specific set of challenges for established firms (and opportunity for start-ups). Importantly, these frameworks concerning the interplay between innovation and competitive dynamics take the nature of the technological trajectory under consideration as given in order to draw out the organizational and competitive implications of that type of innovation. The early work of Henderson and Christensen separately consider the potential uncertainty and path-dependent nature of innovation and the "limits" of the Technology S-curve as a technological construct (Christensen 1992b,c; Henderson 1995). However, the paradigmatic shift in thinking ushered in by their work is grounded in their ability to offer concrete frameworks linking the emergence of specific types of technological trajectories to specific predictions regarding competitive dynamics (Gans 2016).

The concepts of "disruptive" and "architectural" innovation have been among the most influential contributions to our academic and managerial understanding of innovation strategy (Gans 2016). Over the past quarter century, a significant academic and managerial literature has developed that have by and large taken the underlying premises of both theories as given and extended or critiqued them on a theoretical or empirical basis. On the one hand, Henderson and Clark (1990) offered a foundational stepping stone upon which to consider the economic, strategic, and institutional conditions under which established firms are able to manage technological transition, including the potential mitigating impact of complementary assets (Tripsas 1997), the role of managerial cognition (Tripsas and Gavetti 2000), and the potential for strategic cooperation between incumbents and entrants (Gans and Stern 2000). On the other hand, Bower and Christensen (1995) and Christensen (1997) offered a synthetic and accessible overview of the theory of disruptive innovation which had an immediate and persistent impact on managerial practice and thought (Christensen et al 2018; Gans 2016). A significant body of research drew out the economic, institutional and strategic conditions under which this theory applies (Adner 2002; Henderson 2006; Adner and Kapoor 2016), with a separate body of work offering critiques of its logical consistency or empirical relevance (King and Tucci 2002; Danneels 2004; Lepore 2014; King and Baatartogtokh 2015).

At the heart of this ongoing and impactful research agenda are two interrelated concepts: (a) certain types of technological breakthroughs result in new and distinct Technology S-curves and (b) these Technology S-curves can have significant consequences for competitive and organizational dynamics. But, these foundational premises depend importantly on the (often implicit) assumption that the evolution of S-curves is autonomous from innovation strategy itself. For the emergence of an S-curve to have consequences, S-curves themselves must evolve independently from the choices of those firms that the S-curve is supposed to influence. Concretely, in our discussion of start-ups Better Place and Tesla, there was an underlying prediction that a new battery for electric vehicles S-curve would have the potential to achieve a specific performance target (in a given amount of time) to compete against traditional automobile manufacturers. This premise is based on an (implicit) assumption that the evolution of battery performance is a (perhaps difficult to forecast) function of innovative investment in batteries. However, the battery S-curve was not simply a function of raw innovative investment but also a function of strategic choices made by Better Place and Tesla; whereas Better Place took a more aggressive stance that allowed more rapid early progress that ultimately stalled, Tesla undertook a longer-term exploratory approach that has allowed them to take more complete advantage of the underlying potential of lithium ion battery technology.

Put another way, there are two seemingly contradictory conceptions of the Technology S-curve. On the one hand, the origins or the S-curve aimed at analysis at the level of technological trajectories takes the performance profile of the S-curve as given (if uncertain at the outset), and draws out implications for strategy and performance (and is indeed the foundation for the theory of disruptive innovation, which has become ubiquitous among practitioners and entrepreneurs). At the same time, the Technology S-curve is a consequence of innovation strategy at the firm level, most notably reflecting the way in which organizational learning proceeds from an orientation towards exploration to a focus on exploitation.

But, if the S-curve is a technological concept, in what sense does choice of innovation strategy matter? And, if the result of managerial choice, what role does the S-curve play in innovation strategy and competitive dynamics?

III. A Choice Based Approach to the Technology S-curve

To revolve this paradox and clarify the implications technology strategy for start-ups, we reformulate the foundations of the Technology S-curve using a choice-based framework. Rather than taking the Technology S-curve as a purely technological phenomena, the Technology S-curve realized by a given start-up also reflects the strategic choices of the entrepreneur. Concretely, for a given emerging technological trajectory, there exists an envelope of potential innovation outcomes; the precise trajectory that is realized depends crucially not simply on technological potential but on the strategic choice that entrepreneurs and managers make in prioritizing exploration versus exploitation at different stages of the innovation process. This strategic choice is, perhaps, most salient for entrepreneurs where choices to prioritize exploration (and long-term potential innovation outcomes) must be balanced against resource constraints and an incentive to partially resolve uncertainty (through exploitation) to acquire more resources (e.g., funding, team).

The challenges facing an entrepreneur choosing to engage with an emerging technological trajectory throws into sharp relief the subtle distinction between the broad organization learning tradeoff between exploration and exploitation (Levinthal and March 1981; March 1991) and the microfoundations of the Technology S-curve for a given technological trajectory. For any given potential trajectory, entrepreneurs face a choice as to how to balance exploration and exploitation at each moment in time; different choices will result in a different trajectory for that innovation. Manso (2011) articulates this logic clearly in the context of a model in which a manager seeking to motivate innovation over multiple periods chooses between an incentive scheme that provides near-term rewards for "base hits" (and penalizes failure) versus an incentive scheme that has more muted incentives in the short term (and indeed "rewards" short-term failure) but provides higher-powered incentives for a long-term "home run." While Manso focuses primarily on this tradeoff as a challenge for innovation management, consider the compounded incentive challenge faced by an entrepreneur. Though at least equally (and in some cases more) incentivized for a long-term "home run" as the manager, accessing the resources required to achieve the long-term "home run" often necessitates (or highly incentivizes) partially resolving the uncertainty of the long-term "home run" through "base hits." Yet, this focus on repeated near-term "base hits" (and accompanying exploitation) may limit the entrepreneur's ability to achieve the long-term "home run," resulting in a different trajectory for the emerging technology. In particular, where the cost of "ambidexterity" (i.e., managing both the exploitation and the exploration) is sufficiently high, an entrepreneur may choose to focus exclusively on commercializing the 'base hit' technology (and never realize the full potential of the Technology S-curve), or bypass the near-term exploitation in order to realize the full potential of the emerging technology. Indeed, these are the routes that were evident in the choices of Better Place and Tesla.

Specifically, for a given emerging technological trajectory, there exists a strategic choice as to whether to prioritize maximal learning through exploration versus the speed of the transition towards exploitation. The existence of an inverse relationship between the speed at which there is a transition towards exploitation and the overall ability to realize underlying technological potential induces a strategic tradeoff between these alternative approaches. Whereas a short period of exploration allows the start-up to gain a short-term performance benefit, a longer period of exploration allows the start-up to ultimately reach a higher level of technological performance. This insight motivates a reconceptualization of the Technology S-curve as an envelope of potential S-curves, only one of which is realized by the start-up (see Figure 1). In its starkest form, a start-up that prioritizes near-term performance (the red curve) sacrifices their ability to realize long-term potential, while a start-up that is willing to undertake a longer period of exploration (the blue curve) bears the costs of lower performance at the earliest stages of the trajectory though may realize greater long-term potential.

[Insert Figure 1 Here]

Figure 1 captures the essential tradeoff confronting any entrepreneur engaging with an emerging technological trajectory – whether to prioritize activities and investments that lead to higher short-term performance (but lower long-term potential) versus activities and investments that are less likely to yield short-term advantage but may allow for advantage over a longer period of time (Manso 2011). This tradeoff moreover manifests itself in terms of concrete choices facing start-up firms. Perhaps most directly, any entrepreneur considering the potential commercialization path for an emerging technology must choose the incentive system, organizational structure and timeline provided to their technical personnel (which might include the entrepreneur themselves). A more exploitation-oriented approach involves the provision of high-powered incentives for bringing any product "first to market," the integration of technical staff into "product" teams, and the assignment of short-term meaningful deadlines to establish a commercialization "pathway." Conversely, tolerating early-stage failures but providing long-term rewards for "breakthrough" product achievements, a more insular approach towards technical development, and explicitly holding off on exactly how to commercialize the emerging technology are each hallmarks of a more exploration-oriented approach.

Manso and Ederer (2013) provide detailed laboratory-based empirical evidence for the salience of the tradeoff between exploitation versus exploration-oriented incentives on behavior. Specifically, the core experiment allows a comparison of the degree of "search" behavior (in the form of considering alternative locations, recipes and prices for a lemonade stand) by one group whose performance depends on their average level of profitability across the experiment (i.e., a pay-for-performance contract) versus those who are compensated on the profitability achieved by their lemonade stand during only the latter half of the experiment (i.e., an exploration-oriented contract). Manso and Ederer document striking differences in behavior and performance: whereas only 40% of the exploitation-oriented, pay-for-performance group identifies the most profitable location, 80% of those provided exploration-oriented, long-term failure-tolerant incentives are ultimately able to find the best location. Rather than reflecting differences in the nature of the opportunity, this experiment highlights how differences in incentives lead to differences in performance for a given opportunity.

The sharp tradeoff between a short-term and long-term orientation depicted in Figure 1 raises the natural question of how salient this tradeoff might be in practice. For example, it is possible that a start-up that prioritized exploitation at a relatively early stage could simply revert back to a more exploratory mode (perhaps resulting in an ascending wave pattern in which a start-up pursuing the red trajectory could then transition back to the blue trajectory as they exhaust the value of their limited early-stage learning). However, a shift towards exploitation nonetheless involves shifting incentives, personnel, and activities in a way that likely induce at least some commitment away from more exploration-oriented activities. For example, even after the limitations of Better Place's approach to the exploitation of lithium ion batteries became apparent, that firm was unable to then make a transition to join Tesla in a more fundamental innovation process in pursuit of practical electric vehicles. This tradeoff is made even more salient by the fact that the strategic management of an emerging technological trajectory is but one of several interdependent choices made by start-up innovators, with the combination resulting in both tradeoffs and limited commitments (Gans Stern and Wu 2019; Gans Scott and Stern 2020).

A second and perhaps more foundational objection lies in whether it is possible for an entrepreneur to meaningfully forecast the potential for a given technological trajectory at the time when that trajectory is just emerging and the start-up must choose between alternative commitments. The history of technology is littered with examples of the inability and inaccurate forecasts by innovators of the potential for a technology at the moment of its initial development (Rosenberg 1994). And, from a more contemporary perspective, there are sharp debates among researchers and entrepreneurs about the potential for emerging technologies such as artificial intelligence, autonomous vehicles, blockchain technology, or synthetic biology. The choice-based approach redirects attention away from the abstract question of whether a technology has reached its limits towards the question of how strategic choice affects the ability to reach a given level of technological performance.

Finally, it is important to clarify that the relationship between a more exploration-oriented versus exploitation-oriented approach and competitive advantage is ambiguous. While the more exploratory technology strategy of Tesla has ultimately allowed Tesla to establish itself as a global automotive company, there are many instances where the more exploitation-oriented strategy allows a start-up to attract resources and commercialize more quickly.

Consider the classical "innovation" race during the 1970s to leverage the emerging tools of biotechnology to produce human insulin. The two most advanced teams pursuing the most important early-stage product of biotechnology were Genentech (founded by Herb Boyer of UC-SF and Bob Swanson, a venture capitalist) and Biogen (whose founding team included Harvard biologist Wally Gilbert). Though both teams were focused on exploiting the same emerging Technology S-curve (for the same application!), the teams undertook two very different approaches (Hall 1988; Stern 1995). On the one hand, the Biogen team pursued a more exploratory, scientifically interesting "complementary DNA" approach in which human insulin would be produced by first extracting mRNA, transforming that material into human insulin DNA, and then cloning that human DNA material. Though this approach faced significant obstacles (e.g., there was extensive regulation governing the use of human DNA materials), this approach ensured that the team, once successful, would have acquired a systematic scientific understanding of the process by which actual human DNA could be reliably extracted and replicated. In contrast, the Genentech team simplified their pathway to commercialization by the "less interesting" approach of simply building the DNA fragments using commercially available materials. As recounted by Keiichi Itakura, one of the key researchers on the Genentech project, "When we started the insulin project, DNA synthesis was not the risky part." (as quoted in Stern 1995, p. 172). By focusing more intensely on exploitation, Genentech was indeed able to "win" the race to produce synthetic insulin, allowing it to establish a license with Eli Lilly and raise capital through a (highly successful) IPO (Stern 1995).

IV. A Simple Model of Start-up Technologies

Building on this discussion, we now develop a simple (essentially decision-theoretic) model of start-up technology strategy. Specifically, we focus on how a start-up confronting a potential emerging Technology S-curve chooses the timing of when to switch from exploration towards exploitation, and the implications of that endogenous strategic choice for the realized evolution of a particular firm-level technological trajectory (with implications for the industry-level to follow in the next section). Though highly stylized, we are able to clarify the role of technological uncertainty on start-up technology strategy, the impact of constraints (such as limited resources) on technological evolution, and how technology choice is shaped by the possibility of imitation.

The model clarifies the traditional logic of how an entrepreneur pioneering an emerging Technology S-curve can utilize entry on the basis of a more limited exploration S-curve (e.g., the red S-curve) as an experiment on the path towards subsequent performance improvements (and return to the blue S-curve). However, the model also identifies the potential for entrepreneurs themselves to be trapped by a novel form of the "replacement effect" whereby their very success with limited exploration reduces their incentives to explore the full potential of a novel technological trajectory.

A Single Exploration Choice

Consider an entrepreneur who perceives a potential emerging technological trajectory, such as the broad opportunity in lithium ion batteries perceived by Elon Musk and his collaborators in the mid-2000s. The entrepreneur's evaluation of the potential of the emerging trajectory rests on their ability to realize at least some minimum level of technological performance from the emerging technology itself. At the same time, as the distinct strategic choices of the founders of Tesla and Genentech highlight, entrepreneurs need not choose their entrepreneurial strategies in a manner that requires achieving the highest-level of performance from the emerging trajectory. Entrepreneurs, reflecting on their desired realized level of performance from the technology, are making a strategic choice as to the level of investment in exploration in the emerging Technology S-curve. To focus on the interplay between the perception of opportunity (Shane 2000) and entrepreneurial choice (Gans Stern and Wu 2019), we begin with the assumption that the choice of the entrepreneur shapes the evolution of their firm-level Technology S-curve and model a start-up's strategic choice over the investment in exploration (and performance) of an emerging Technology S-curve. In essence, the entrepreneur is facing a strategic choice over the set of potential S-curves that may be realized based on differing levels of investment in exploration.

It is useful to begin with a simple characterization of the potential technological trajectory that the entrepreneur can explore and ultimately develop. Specifically, consider a standard functional form for the S-curve which depends on the overall technological potential for the S-curve (K^{MAX}), the length of time that is spent in exploration before releasing a product (t), the amount of exploration before an inflection point is reached (T) and the steepness of the S-curve along that path (k):

$$\kappa(t) = \frac{K^{MAX}}{1 + e^{-k(t-T)}}$$

For simplicity, we consider the case where we take the limit as $k \to \infty$, and so (1) becomes a step function where the value of $\kappa(t) = 0$ for $t < T$ and $\kappa(t) = K^{MAX}$ for $t \geq T$. Thus, for a given S-curve, the entrepreneur will want to choose a length of time, T, to spend in exploration.

Now consider, the entrepreneur's choice of S-curve itself. The elements of a choice set of S-curves with individual elements labelled i, are described by pairs $\{K_i^{MAX}, T_i\}$. We assume that the discount factor for the entrepreneur is $\delta \in [0,1)$ and, thus, for a given K_i^{MAX}, the entrepreneur would prefer a shorter T_i and, for a given T_i, the entrepreneur will prefer a higher K_i^{MAX}. Thus, what is relevant is the upper envelope of the set of S-curves. To make this transparent, we assume that $K^{MAX} = T$ so we can think of the entrepreneur as selecting an S-curve on the basis of the length of time spent in exploration. That time, T, delays when the product can be exploited but also increases the potential performance associated with that product. We will assume that performance drives the value of the start-up's product to its consumers with higher performance being synonymous with greater value; that is, the expected profits of the venture, Π, are increasing in T.

As a baseline, suppose that there is there is no uncertainty and the decision to invest in exploration is a one-time choice (i.e., there is a single period of exploration followed by commercialization), then the entrepreneur chooses to solve:

$$Max_T \frac{\Pi(T)}{(1+\delta)^T} - c(T)$$

T^* is determined by $\frac{1}{(1+\delta)^{T^*}}\left(\Pi'(T^*) - \log(1+\delta)\,\Pi(T^*)\right) = c'(T^*)$; so a venture that is more patient will choose a higher *T*. Moreover, if the start-up is resource constrained with total resources, $\underline{c} = c(T^C) < c(T^*)$, the chosen S-curve resolves to positive performance in a shorter amount of time but with a lower potential.

In reality, of course, there is the potential for uncertainty with regard to the success of exploration, given by a probability $p(T)$, as well as whether exploitation will result in imitation, that is, the public nature of exploitation leads to a probability, q, that the start-up will face perfect competition ex-post. We assume that $p(T)$ is non-decreasing implying that more uncertainty is resolved if there is commitment to a longer time in exploration. The probability of imitation, q, is assumed to be invariant with T as there is no clear theoretical rationale as to why it would increase or decrease with the amount of exploration as the result is a clearly implementable idea. Given this, we write the start-up's choice problem as:

$$Max_T \frac{p(T)(1-q)\Pi(T)}{(1+\delta)^T} - c(T)$$

Note that the optimal T is now decreasing in q, the probability that the start-up will face perfect competition ex-post, and that the start-up considers the reduction in uncertainty regarding exploration success as a positive factor in terms of choosing a longer exploratory period.

Note that q can capture various aspects of competition. It could be pure imitation that might arise if the new technology is freely available but does not have, say, intellectual property protection. It could also represent the probability that an incumbent responds quickly or not. This gives a window on potential disruptive forces. For instance, the start-up might otherwise choose to enter a market targeting niche customers rather than an incumbent's existing customers. If they do this, then q will be correspondingly lower and so the start-up would choose a longer period of exploration. In this respect, we can see how the start-up's technology choice will interact with its other strategic choices (Gans Scott and Stern 2020).

A Repeated Exploration Choice

Thus far, we have modelled the start-up's technology choice as a single decision regarding T. In reality, start-ups might reset or alter their trade-offs overtime. For instance, they may initially choose a quick exploration cycle in order to resolve uncertainty sooner (and better acquire resources and key learnings) and keep the option of further exploration and an elongated S-curve for the future. The most straightforward mechanism limiting this type of "surfing the frontier" would be if there exists a tradeoff between the commercialization of the technology at one performance level and unfettered exploration of that same trajectory in order to achieve an even higher level of performance. Thus, it is interesting to analyze whether there is any advantage to the start-up making up-front commitments to exploration rather than preserving options to invest later on.

To model this, we suppose that exploitation now has an explicit cost, C, for each time a new product is introduced. Recall that, we suppose that exploration resolves uncertainty $p(T)$. Thus, it is now possible that a start-up might choose to explore for a shorter period of time in order to exploit and yield an earlier signal on the technology's potential success. In particular, having exploited, the start-up either finds out, with probability $1 - p(T)$, that no further exploration will be useful or with probability, $p(T)$ that further exploration will yield the benefit, T with certainty. This can allow it to save on exploration costs and it also may be a mechanism of securing additional funding and resources. That said, the publicness of exploitation implies that there is still a risk, at any point, that the product is imitated. Nonetheless, even with imitation of a successful product, the start-up can still choose to embark on additional exploration.

Given this, the start-up now chooses T to solve:

$$Max_T \frac{p(T)\big((1-q)\mathrm{V}(T) + qv(T)\big) - C}{(1+\delta)^T} - c(T)$$

where $\mathrm{V}(T) = \max_{T'} \frac{\Pi(T)}{(1+\delta)^T} + \frac{(1-q)\Pi(T') - \Pi(T) - C}{(1+\delta)^{T'}} - c(T')$ (if there is no imitation of the initial product) and $v(T) = \max_{T'} \frac{(1-q)\Pi(T') - C}{(1+\delta)^{T'}} - c(T')$ (if there is imitation of the initial product). Each of these are the option value of continuing exploration after you have exploited (that is, bumping up to a better S-curve).

The most significant finding here is that, where there is no imitation, the optimal *T*' is increasing in *T*, the initial level of exploration. That is, the higher the S-curve potential reached in the initial round of exploration, the higher is the target for an S-curve in subsequent rounds. Intuitively, if the initial S-curve has positive value when exploited that signals that further investment in S-curves will also have positive value. So, one effect, as expected, is to spur additional investment following the initial experiment. However, importantly, the threshold to justify that additional investment must be correspondingly higher because that investment will replace the initial successful investment. In other words, there is a *replacement effect* that emerges in the second stage but in contrast to that discussed by Arrow (1962), this applies to start-ups. Ironically, if there is imitation, this start-up replacement effect is not present and the optimal length of exploration is shorter.

This gives rise to an interesting insight. Being able to explore for a time and then exploit, changes the potential behavior of the start-up. Specifically, if the start-up were to have a goal of a long total period of exploration in order to end up on an S-curve with higher potential, to exercise the option associated with staged exploration, the start-up would want to choose a more limited period of exploration initially in order to *commit* themselves to continuing to that higher level should the technological uncertainty resolve favorably. Whether the technological choice ends up being an S-curve with greater exploration and greater potential or not depends upon the trade-off between a desire for an early resolution of uncertainty versus a desire to economize on resource constraints by committing to a longer period of exploration. This “commitment” approach to long-term S-curve exploration is indicative of the commitment value one sees in Tesla.

The interaction of the entrepreneur’s technological choice with their choice of beachhead market illustrates how lacking a pre-commitment to exploration, an entrepreneur may ultimately choose not to pursue a S-curve with higher potential despite earlier intentions. For instance, following Christensen (1997), we might think of start-ups who are conducting an experiment as choosing a more modest S-curve with which to penetrate an under-served market niche (who place a value of lower technical performance). They could enter that niche successfully (and not encourage imitation) but find themselves, precisely for that reason, likely to go no further and enter mainstream markets because choosing an S-curve with high performance at that point will cannibalize their own position in the niche market as well as potentially provoke a response from established firms (see Gans 2016). In other words, a start-up pursuing an experiment might find themselves stuck incentive-wise from pursuing a technology to a much higher level and scale.

This potential for being stuck at a lower performance level is likely if the start-up is resourced constrained. A start-up that is resourced constrained, in terms of $\underline{c}$ that limits the initial exploratory choice (T) both increases the risk that exploration will not be successful and causes it to anticipate choosing a lower T' should it be successful and further funding is available. This suggests that a more complete analysis of resource constraints, funding commitments and the role of competition and imitation are a fruitful area for future analysis of what drives start-up growth.

V. Industry-Level Technology S-Curve

We have focused our analysis so far on the case of a single start-up firm that is able to exercise significant level of choice over both whether a particular trajectory is developed at all, and, if so, the balance between exploration and exploitation that yields a particular S-curve for that technology. Of course, most emerging technological opportunities – such as electric vehicles – attract multiple potential entrants including existing firms. Indeed, while its micro-foundations are at the firm level, the Technology S-curve is often taken to reflect an industrywide (or sectoral) phenomena (Abernathy and Utterback 1978; Anderson and Tushman 1990; Utterback 1994), and, though often invoked as a strategic tool for individual firms, both the existence and nature of the Technology S-curve are largely grounded in industry-level case studies (Foster 1986; Christensen 1997).

The linkage between the tradeoffs impacting a single firm and the patterns that emerge at an industry level results from an (often implicit) set of assumptions about how a group of firms responds to the emergence of a technological opportunity. Specifically, as illustrated in the rich coevolutionary models of innovation and competition such as Utterback and Abernathy (1975) and Klepper (1996), the key dynamic is the shared challenges among firms at different stages of the technology lifecycle result in distinctive patterns of innovation at the industry level. Specifically, in the early stages after a technological discontinuity, a wide range of firms enter and engage in experimentation and learning – this period corresponds to the initial flat portion of the S-curve. However, when technical and institutional factors allow for the establishment of a dominant design, there is a significant uptick in the realized performance of the technology (expanding demand), a shakeout of firms, and a shift towards more incremental process innovation, in a way that mirrors the pattern of increasing and then declining returns to R&D effort.

The choice-based perspective developed earlier provides insight into both the underlying assumptions of such a model, as well as insight into the nature of the strategic choices facing start-up innovators when multiple firms are pursuing an emerging technological opportunity. Notably, connecting the tradeoff between exploration and exploitation at the firm and industry level requires implicitly at least three assumptions in evolutionary models: First, that there is a shared perception of the opportunity presented from the emerging technological trajectory. Second, that there is shared access to the emerging trajectory itself, with few restrictions on entry due to intellectual property or secrecy. Third, that there is shared learning as firms engage with the emerging trajectory; that the experiments and failures from individual firm's exploration activities are (at least partially) transmitted to other firms within the industry and contribute to accumulated learning at the industry level. Though industry ferment is characterized by broad experimentation and competing paths (Clark 1985), it relies on a degree of common perception, access, and learning across entrants (both entrepreneurs and existing firms) engaging with the emerging trajectory.

Perhaps more saliently, a choice-based perspective offers novel insight into the core strategic choices facing firms engaging with an emerging technological trajectory. Rather than a "passive" process by which firms engage first in exploration and then shift towards exploitation with the emergence of a dominant design, a choice-based approach suggests that firms face a core strategic tradeoff between attempting to shift from exploration to exploitation at the earliest stage possible (in order to gain competitive advantage through the establishment of the dominant design) versus continuing to invest in exploration in order to enhance the likelihood that a shift towards exploitation will in fact be successful. In other words, building on the model in Section IV, each firm in the emerging trajectory faces a tradeoff between exploration and exploitation, but there is significant uncertainty for each firm about which mix will yield a level of learning and subsequent technological performance that will be sufficient to catalyze standardization around a dominant design. Specifically, a firm that shifts from exploration to exploitation prematurely has a higher chance for failure; a firm that maintains an orientation towards exploration risks missing the shift in industry dynamics. Ultimately, this choice-based perspective suggest that firms are not simply choosing among technological S-curves with which to engage, but also choosing how to shape the S-curve of the technologies they engage.

This broader lesson on dynamics offers insights about the emergence of Tesla over the past 15 years as the leading firms in electric vehicles. Tesla was neither the first start-up to enter the electric vehicle industry, nor was it the only firm that sought to take advantage of the potential for lowered cost and increased power from lithium ion batteries. However, in contrast to both efforts by existing incumbent firms in the combustion engine automobile, and other electric vehicle start-ups, Tesla both committed itself to a meaningful level of exploration that allowed it to actually realize a higher level of technological performance and so land a beachhead market, and also committed itself to experimental commercialization that avoided more exploratory efforts (such as relying on new electrolytes or load-bearing battery cells) that have yielded significant learning but where the firm never shifted towards meaningful exploitation and commercialization.

VI. Implications for Strategy and Policy

A choice-based approach to the Technology S-curve offers insights for researchers, practitioners and policy. On the one hand, an immediate consequence of a choice-based approach is that the incentives to pursue exploration versus exploitation around a given technology depend not only on the nature of technology itself, but on the strategic and institutional environment in which that technology is situated. A trajectory being pursued by a long-term oriented entrepreneur (and similarly patient investors and funders) committed to cumulative knowledge development is likely to have few short-term successes but may serve as a foundation for long-term innovation potential. Beyond the laboratory-based evidence of Ederer and Manso (2013), Tian and Wang (2014) and Nanda and Rhodes-Kropf (2013) demonstrate that the failure tolerance and risk appetite of venture capitalists allows for more exploration which ultimately yields greater commercialization of innovative technologies. Alternatively, short-termism in innovation can not only limit the immediate application of a technology but may reduce the long-term potential of that trajectory through a lock-in to a more exploitation-oriented approach.

One of the novel implications of our analysis is that, even when there exists an option to shift back and forth between exploration and exploitation, an early focus on exploitation may induce a “replacement” effect that diminishes incentives to shift back towards a more exploratory mode. Concretely, though the disincentives for innovation arising from the replacement effect have traditionally been linked to incumbent players, the potential existence of an early-stage exploitation-oriented market of an emerging technology may limit the long-term evolution of that path (at least by that start-up firm). Put another way, our model highlights the importance of strategic commitment in being able to realize a more exploratory approach. This potential for lock-in has implications not only for entrepreneurs but also for investors and other stakeholders. Specifically, unless funders are able to commit to a potential staging of funding that commits the start-up to a higher level of exploration, lack of resources is likely to induce start-ups to undertake a more exploitation-oriented approach.

This dynamic is particularly salient in assessing the interaction between a start-up's initial choice of customer, or beachhead market, their technology choice, and their overall entrepreneurial strategy. For instance, an entrepreneur following a disruption strategy (Christensen 1997) by initially pursuing a poorly served niche market (who place a value of lower technical performance) on a strategic path to enter the mainstream market may find that successful entry into the niche market (without threat of imitation) may mute their incentive to continue to the mainstream market (due to cannibalization and potential incumbent competitive response). More generally, the ex-ante strategic freedom entrepreneurs possess that enables them to take advantage of emerging Technology S-curves (while existing firms delay) may be short-lived. Early-stage strategic choices induce organizational inertia (Kapoor and Furr 2015), foreclose particular options (Gans Stern and Wu 2019) and impose subsequent biases (Kaplan and Tripsas 2008).

Though we have focused our analysis on the broad tradeoff between exploration versus exploitation, a similar dynamic emerges if one considers related dimensions of technological change. Specifically, though start-ups may face significant latitude in choosing between a more modular versus more systemic "version" of a technology, choosing the initial degree of modularity within a design constrains design freedom going forward (Egan 2013). For example, the choice by Steve Wozniak and Steve Jobs to develop a more integrated design in the Apple II relative to more modular approaches such as the IBM PC induced a design contrast that has persisted for more than forty years. Similarly, entrepreneurs pursuing an emerging technological trajectory often have significant choice over the degree of generality of that technology (i.e., whether development focuses on specific applications or involve a more general-purpose approach). For example, after a regulatory shock raised the cost of integrating into more specialized markets, firms in the laser industry pursued more "generalist" approaches to that technological trajectory (Conti et al 2019).

More generally, while the bulk of theoretical and empirical research in innovation strategy takes these two dimensions of technology as givens (Bresnahan and Trajtenberg 1995; Schilling 2000), a choice-based approach suggests that each of these aspects of a technology are endogenously influenced. Consider again the contrasting development approaches of Better Place and Tesla. While both companies were premised on taking advantage of the potential for improvement in lithium ion batteries, each company took a very different approach in how to take advantage of that possibility. On the one hand, Better Place focused relatively rapidly on implementing their "swappable battery" concept, resulting in two key design choices. First, to implement a "plug-and-play" automobile battery, Better Place found that it had to fundamentally redesign the automobile itself (most specifically, it had to reconceptualize the traditional drive train upon which the bulk of automobile system design is built). Second, to accelerate time to market (i.e., their focus on rapid exploitation), Better Place chose to undertake its development in a manner that was specifically optimized for automotive applications rather than the broader battery market. Though facing the same opportunity, technology and market, the approach by Tesla during the second half of the 2000s differed sharply on both of these dimensions. The early years of Tesla were marked on the one hand by a design focus that allowed Tesla to take advantage of existing automobile system design (i.e., preserving the traditional drivetrain) while also focusing their technology development in a way that allowed them to lead battery development not simply for automobiles but for a wider range of applications (Ramey 2017).

Finally, our analysis offers significant implication for innovation policy. First, while our analysis has focused primarily on the technology strategy choices of start-up firms, the broad tradeoff we analyze also is salient for individual scientific researchers. As documented in Azoulay, et al (2011), the provision of more exploratory research incentives for life science researchers (resulting from being granted a Howard Hughes Medical Investigator grant) results in a similar quantity of scientific output (i.e., raw numbers of papers) but is associated with the pursuit of more exploratory research approaches (e.g., a higher fraction of papers which either fail (e.g., are poorly cited) or become home runs). Our analysis is moreover consistent with a broader body of evidence (drawn from detailed analyses of disparate publicly funded start-up innovation programs (including DARPA, ARPA-E, NIH, and SBIR) that a central dimension of innovation policy is not simply the provision of research funding but the management and incentives governing the research that is funded (Azoulay et al 2019). More generally, our analysis suggests that, given that, almost by definition, we do not observe the full range of potential ways that any particular technology could be developed, and so it is worthwhile to explicitly consider how entrepreneurs and other decisionmakers make choices about which paths to explore and which are left behind.

References

Abernathy W (1976) Production process structure and technological change. *Decision Science.* 7(4):607-619.

Abernathy W, Utterback J (1978) Patterns of industrial innovation. *Technology Review.* 80(7):41-47.

Adner R (2002) When are technologies disruptive? A demand-based view of the emergence of competition. *Strategic Management Journal*. 23(8):667–88.

Adner R, Kapoor R (2016) Innovation ecosystems and the pace of substitution: Re-examining technology S-curves. *Strategic Management Journal*. 37(4): 625-648.

Anderson P, Tushman M (1990) Technological discontinuities and dominant designs: A cyclical model of technological change. *Administrative Science Quarterly*. 35(1990):604-633.

Arrow K (1962) Economic welfare and the allocation of resources for invention. Nelson R, ed. *The Rate and Direction of Inventive Activity: Economic and Social Factors* (National Bureau of Economic Research), 609-626.

Azoulay P, Fuchs E, Goldstein A, Kearney M (2019) Funding breakthrough research: promises and challenges of the "ARPA Model." Lerner J and Stern S*,* eds. *Innovation Policy and the Economy*, Volume 19.

Azoulay P, Graff Zivin J, and Manso G (2011) Incentives and creativity: evidence from the academic life sciences. *RAND Journal of Economics.* 42(3):527-554.

Baker D (2016) Elon Musk: Tesla was founded on 2 false ideas, and survived anyway. *SFGate (May 31 2016).*

Bower JL, Christensen C (1995) Disruptive Technologies: Catching the Wave. *Harvard Business Review*. 73(1):43–53.

Bresnahan T, Trajtenberg M (1995) General purpose technologies 'engines of growth'? *Journal of Econometrics*. 65(1):83-108.

Burns T, Stalker, GM (1961) *The Management of Innovation* (Oxford University Press).

Carroll GR, Bigelow LS, Seidel ML, Tsai LB (1996) The fates of de novo and de alio producers in the American automobile industry 1885-1981. *Strategic Management Journal.* 17:117-137.

Chafkin M (2014) A broken place: The spectacular failure of the startup that was going to change the word. *The Atlantic (April 2014)***.**

Christensen C (1992a) Doctoral Dissertation, Harvard Business School, Cambridge, MA.

Christensen C (1992b) Exploring the limits of the technology S-curve, Part 1: Component Technologies. *Production and Operations Management.* 1(4):334–357.

Christensen C (1992c) Exploring the limits of the technology S-curve, Part 2: Architectural Technologies. *Production and Operations Management.* 1(4):358–366.

Christensen C (1993) The rigid disk drive industry: A history of commercial and technological turbulence. *The Business History Review*. 67(4):531-588

Christensen C (1997) *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail* (Harvard Business School Press, Boston, MA).

Christensen C, McDonald R, Altman EJ, Palmer J (2018) Disruptive innovation: An intellectual history and directions for future research. *Journal of Management Studies.* 55(7):1043-1078.

Clark KB (1985) The interaction of design hierarchies and market concepts in technological evolution. *Research Policy*. 14(5): 235-251.

Conti R, Gambardella A, Novelli E (2019) Specializing in general purpose technologies as a firm long-term strategy. *Industrial and Corporate Change.* 28(2):351-364.

Danneels E (2004) Disruptive technology reconsidered: A critique and research agenda. *Journal of Product Innovation Management*. 21(4):246–58.

Dosi G (1982) Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change. *Research Policy*. 11(3):147-162.

Ederer F, Manso G (2013). Is pay for performance detrimental to innovation? *Management Science*. 59(7):1496-1513.

Egan E (2013) How start-up firms innovate: technology strategy, commercialization strategy, and their relationship. Working Paper. Available at SSRN: http://ssrn.com/abstract=2364096.

Foster RN (1986) *Innovation: The Attacker's Advantage* (Summit Books, New York, NY).

Furr N, Cavarretta F, Garg S (2012) Who changes course? The role of domain knowledge and novel framing in making technology changes. *Strategic Entrepreneurship Journal.* 6(3): 236-256.

Ganco M, Agarwal R (2009) Performance differentials between diversifying entrants and entrepreneurial start-ups: A complexity approach. *Academy of Management Review.* 34(2): 228-252.

Gans JS (2016) *The Disruption Dilemma* (MIT Press, Cambridge, MA).

Gans JS, Scott EL, Stern S (2018) Strategy for start-ups. *Harvard Business Review*. 96(3):46-51.

Gans JS, Scott EL, Stern S (2020) *Entrepreneurial Strategy.* Draft Book Manuscript.

Gans JS, Stern S (2000) Incumbency and R&D incentives: Licensing the gale of creative destruction. *Journal of Economics and Management Strategy.* 9(4):485-511.

Gans JS, Stern S (2003) The product market and the market for 'ideas': commercialization strategies for technology entrepreneurs. *Research Policy*. 32(2):333-350.

Gans JS, Stern S, Wu J (2019) Foundations of entrepreneurial strategy. *Strategic Management Journal.* 40(5):736-756.

Hall SS (1988) *Invisible Frontiers: The Race to Synthesize a Human Gene* (Sidgwick & Jackson, London).

Henderson R (1988) The Failure of Established Firms in the Face of Technological Change: A Study of Photolithographic Alignment Equipment. Doctoral Dissertation. Harvard Business School, Cambridge, MA.

Henderson R (1995) Of life cycles real and imaginary: The unexpectedly long old age of optical lithography. *Research Policy*. 24(4): 631-643.

Henderson R (2006) The innovator's dilemma as a problem of organizational competence. *Journal of Product Innovation Management*. 23(1):5-11.

Henderson R, Clark K (1990) Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly.* 35(1):9–30.

Kaplan S, Tripsas M (2008) Thinking about technology: Applying a cognitive lens to technical change. *Research Policy.* 37(5):790-805.

Kapoor, R, Furr NR (2015) Complementarities and competition: Unpacking the drivers of entrants' technology choices in the solar photovoltaic industry. *Strategic Management Journal.* 36(3), 416-436.

Khessina OM, Carroll GR (2008) Product demography of de novo and de alio firms in the optical disk drive industry, 1983-1999. *Organization Science.* 19(1): 25-38.

King AA, Baatartogtokh B (2015) How useful is the theory of disruptive innovation? *MIT Sloan Management Review*. 57(1):77–90.

King AA, Tucci CL (2002). Incumbent entry into new market niches: The role of experience and managerial choice in the creation of dynamic capabilities. *Management Science*. 48(2):171–86.

Klepper S (1996) Entry, exit, growth and innovation over the product life cycle. *American Economic Review*. 86(3):562-83.

Klepper S, Simons KL (2000) Dominance by birthright: Entry of prior radio producers and competitive ramifications in the US Television receiver industry. *Strategic Management Journal.* 21: 997-1016.

Lepore J (2014) The disruption machine: What the gospel of innovation gets wrong. *The New Yorker (June 23 2014).*

Levinthal D, March, J (1981) A model of adaptive organizational search. *J. Economic Behavior & Organizations*. 2(4):307–333.

Levinthal D, March, J (1993) The myopia of learning. *Strategic Management Journal.* 14(2):95-112.

Manso G (2011) Motivating innovation. *Journal of Finance.* 66(5):1823-1860.

March J (1991) Exploration and exploitation in organization learning. *Organization Science*, 2(1):71–87.

Mitchell W (1994) The dynamics of evolving markets: The effects of business sales and age on dissolutions and divestitures. *Administrative Science Quarterly,* 39(4):575-602.

Nagy B, Farmer JD, Bui WM, Trancik JE (2013) Statistical basis for predicting technological progress. *PLoS ONE.* 8(2): e52669.

Nanda R, Rhodes-Kropf M (2013) Investment cycles and startup innovation. *Journal of Financial Economics.* 110(2)403-418.

Nelson R (1962) The link between science and invention: the case of the transistor. Nelson R, ed. *The Rate and Direction of Inventive Activity: Economic and Social Factors* (National Bureau of Economic Research), 549-584.

Ramey J (2017) The first tesla roadster: A look back at the early adopter's electric car. *Autoweek* (November 27 2107).

Roberts E (1988) What we've learned: Managing invention and innovation. *Research-Technology Management.* 31(1):1988.

Rosenberg N (1976) *Perspectives on Technology* (Cambridge University Press, Cambridge, UK).

Rosenberg N (1994) *Exploring the Black Box: Technology, Economics, and History* (Cambridge University Press, Cambridge, UK).

Sahal D (1981) Alternative conceptions of technology. *Research Policy.* 10(1981):2-24.

Schilling MA (2000). Toward a general modular systems theory and its application to interfirm product modularity. *Academy of Management Review*. 25(2):312–334.

Schumpeter J (1942) *Capitalism, Socialism, and Democracy* (Harper & Row, New York, NY).

Shane S (2000) Prior knowledge and the discovery of Entrepreneurial Opportunities. *Organization Science.* 11(4):448-469.

Stern S (1995) Incentives and Focus in University and Industrial Research: The Case of Synthetic Insulin. *Rosenberg N, Gelijns AC, Dawkins H,* eds. *Sources of Medical Technology Universities and Industry* (National Academies Press, Washington DC), 157-187.

Squatriglia C (2009) Renault's EV features a swappable battery. *Wired* (June 22 2019).

Teece DJ (1986) Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*. 15(6):285-305.

Tian X, Wang TY (2014) Tolerance for failure and corporate innovation. *Review of Financial Studies.* 27(1):211-255.

Tripsas M (1997) Unraveling the process of creative destruction: Complementary assets and incumbent survival in the typesetter industry. *Strategic Management Journal.* 18(1):119-142.

Tripsas M, Gavetti G (2000) Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal.* 21(10/11):1147-1161.

Utterback J, Abernathy W (1975) A dynamic model of process and product innovation*. OMEGA, The Int. Jl of Mgmt Sci*. 3(6): 639-656.

Utterback J (1994) *Mastering the Dynamics of Innovation: How Companies can Seize Opportunities in the Face of Technological Change* (Harvard Business School Press, Boston, MA).

Wu B, Wan Z, Levinthal DA (2014) Complementary assets as pipes and prisms: Innovation incentives and trajectory choices. *Strategic Management Journal*. 35(9), 1257-1278.

Figures

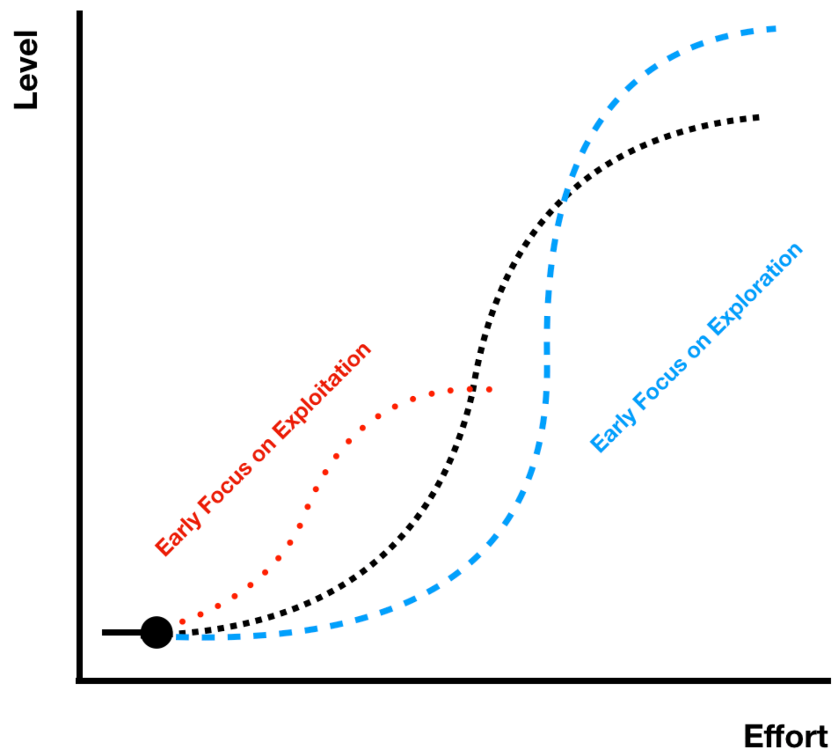

Figure 1: For an emerging technology with which a start-up chooses to engage, the entrepreneur is facing a strategic choice over the set of potential firm-level S-curves that may be realized based on differing levels of investment in exploration (versus exploitation). A start-up that prioritizes near-term performance (the red dotted curve, left) sacrifices their ability to realize long-term potential, while a start-up that is willing to undertake a longer period of exploration (the blue dashed curve, right) bears the costs of lower performance at the earliest stages of the trajectory though may realize greater long-term potential.