NBER WORKING PAPER SERIES

OF ACADEMICS AND CREATIVE DESTRUCTION:
STARTUP ADVANTAGE IN THE PROCESS OF INNOVATION

Julian Kolev
Alexis Haughey
Fiona Murray
Scott Stern

Working Paper 30362
http://www.nber.org/papers/w30362

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2022

This paper was prepared for the Economics of Creative Destruction Festschrift Conference in Honor of Phillipe Aghion and Peter Howitt. The authors thank Phillipe and Peter for their inspiration, and Ufuk Akcigit, John van Reenen, and other conference participants for helpful comments and suggestions. We would also like to thank Mercedes Delgado for providing access to the patent data, and the MIT Technology Licensing Office for access to statistics from the AUTM data. All errors remain our own. The views expressed are those of the individual authors and do not necessarily reflect official positions of the Office of the Chief Economist or the U. S. Patent and Trademark Office. Stern and Murray have provided more detailed financial disclosures on the NBER website. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w30362.ack

Of Academics and Creative Destruction: Startup Advantage in the Process of Innovation
Julian Kolev, Alexis Haughey, Fiona Murray, and Scott Stern
NBER Working Paper No. 30362
August 2022
JEL No. L24,L26,M13,O31,O32,O34

ABSTRACT

What is the role of startups within the innovation ecosystem? Since 2000, startups have grown in their share of commercializing research from top U.S. universities; however, prior work has little to say on the particular advantages of startup ventures in the innovation process relative to more traditional alternatives such as academia and established private-sector incumbents. We develop a simple model of startup advantage based on private information held by the initial inventor, and generate predictions related to the value and impact of startup innovation. We then explore these predictions using patents granted within the regional ecosystems of top-25 research universities from 2000 to 2015. Our results show a significant startup advantage in terms of forward citations and outlier-patent rates. Further, startup innovation is both more original and more general than innovation by incumbent firms. Moreover, startups that survive to become "scale-ups" quickly grow to dominate their regional innovation ecosystems. Our findings have important implications for innovation policy.

Julian Kolev
US Patent and Trademark Office
600 Dulany St
MDE 2D32
Alexandria, VA 22314
julian.kolev@uspto.gov

Alexis Haughey
MIT Sloan School of Management
100 Main St
Cambridge, MA 02142-1508
haugheya@mit.edu

Fiona Murray
MIT Sloan School of Management
100 Main Street, E62-470
Cambridge, MA 02142
and NBER
fmurray@mit.edu

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

Introduction

Innovation serves as an essential engine of modern economies, most prominently as a driver of long-term productivity and growth. Economists have long sought to understand the drivers of innovation as a window into the differential growth rates of nations, regions, and firms. The seminal contributions of our colleagues Phillipe Aghion and Peter Howitt provide foundational insight into the drivers of innovation, and in particular the role that incentives and competition play in shaping innovative investment and effort by different organizations.

The first economic question that arises when considering the rate of innovation is whether (and which) organizations have adequate incentives to undertake the costly investments needed to drive innovation. In "A Model of Growth through Creative Destruction," Aghion and Howitt (1992, hereafter AH) provide the critical building blocks necessary for addressing this challenge by clarifying precisely how firm-level incentives for investment are shaped by the prospect of creative destruction. Specifically, AH highlight the fact that incentives for innovation whose outcome offers a superior alternative to that currently available in the market reside with entrepreneurs rather than incumbent firms (due to the "replacement effect"), but these incentives are themselves shaped by the prospect of subsequent follow-on innovation that will in turn overtake the (transitory) market power of a successful startup. In other words, AH offer a general equilibrium framework in which (a) startup firms play a distinctive and outsized economic role in the process of economic growth, and (b) innovation incentives for new entrants reflect a balance between potential market power for a single period versus the degree to which incentives are provided for subsequent entrants to innovate in future periods.

While AH provide critical insight into the incentives of startup ventures to engage in the process of innovation, their framework abstracts away from the incentives provided to the agents undertaking that research. This is the question that is addressed in the (complementary) theoretical framework of Aghion, Dewatripont and Stein's paper, "Academic Freedom, Private-Sector Focus, and the Process of Innovation" (2008, hereafter ADS). Similar to AH, ADS consider a multi-period model of innovation where incentives for a company to innovate in a given period are shaped by expected returns in some future period. However, relative to AH, ADS emphasize that a successful innovation may involve a long research line with multiple stages, and examine the challenges facing an organization contracting with researchers at each stage of this line. A critical insight from ADS is that, when researchers have a preference to maintain a degree of control over project choices, it is optimal to cede the control over early-stage projects to researchers (similar to the incentives of academia); by contrast, private funders (i.e. investors) will optimally assert control over project direction as the innovation gets closer to the marketplace.

A natural question that arises is how the two perspectives of ADS and AH are related to one another. On the one hand, AH suggest that startups play a particularly important role in driving creative destruction. On the other hand, ADS suggest that a for-profit firm (whether incumbent or startup) is likely to be involved in the late- (but perhaps not early-) stages of commercialization of a multi-stage innovation process. However, neither model directly addresses how the process of creative destruction (in the market) is related to the incentives (within organizations) required to invest in a multi-stage innovation process. The purpose of this paper is to explore the connection between these two perspectives.

Our analysis suggests that startups play a special role in the process of commercializing innovations, linking early-stage scientific research to the broader economy-wide impact of creative destruction. This linkage is particularly important given the increasing reliance on science-driven startups in addressing a range of pressing societal challenges. In so doing, our analysis reflects and extends the durable and far-reaching implications of the Aghion and Howitt creative destruction framework for understanding the relationship among innovation, institutions, and economic growth.

Building on a more in-depth discussion contrasting AH and ADS, our analysis begins with a set of short case studies illuminating the process by which foundational scientific research is transferred from academia into the private sector. In each case, though there is no requirement that a startup venture be involved in the process of commercialization, we illustrate that a critical step in the hand-off from academia to industry takes place through entrepreneurship and startup venture formation. This motivates our more formal analysis, where we extend the ADS model to specify the conditions under which startup ventures will optimally bridge the gap between universities and large firms. In doing so, we identify two distinct drivers of "startup advantage": first, academic startup founders may have access to private information about the value of specific high-risk research lines. Startups are, therefore, willing to pursue these ideas at an earlier stage in the innovation process described in the ADS model. Second, startups do not experience the cannibalization effect described in AH, whereby successful innovation by an incumbent comes at the expense of their own pre-existing competitive advantage. Because of this difference, startups are in a position to pursue more disruptive research lines in which incumbents would normally under-invest or lack interest. Taken together, these two channels form the foundation of "startup advantage," whereby new ventures are uniquely positioned to develop and extend research lines that are distinct from the innovations targeted by incumbents.

Our theoretical analysis generates two central hypotheses. First, startup innovation will be more valuable and ultimately more impactful than that of *either* universities or large firms, and second, startups will generate innovations that are more radical and disruptive than those of incumbent firms. We provide descriptive statistics consistent with these hypotheses using a sample of patents generated in the vicinity of the top 25 research universities in the United States from 2000 to 2015.[1] We find strong evidence for startup advantage in both average forward citations and the rate of outlier patents (in the top 5% of the citation distribution), supporting our first hypothesis. We also find that startup patents score higher in terms of originality and generality relative to patents from established firms (but not significantly different from university patents), which is consistent with our second hypothesis. Finally, we show direct evidence for creative destruction by tracking the development of startups over time: they play an increasingly prominent role in their regional innovation ecosystems. Overall, our findings suggest that startup innovation is qualitatively different from the innovation in other organizational settings: there is a clear "startup advantage" in the quality and impact of startup patents relative to established firms.

Related Literature and Motivation

A Tale of Two Theoretical Frameworks

The foundation of our analysis is grounded in endogenous growth theory, where resources dedicated to research and innovation generate improvements in productivity and thus in output quality. In AH, firms compete to develop an innovation which consists of a new intermediate good. This new intermediate good results in more efficient production of a consumption good, in turn driving growth in the economy. The innovating firm is able to attain an immediate (transitory) monopoly until the next innovation occurs. In this framework, there is a natural tendency towards creative destruction, driven by the Arrow replacement effect (Arrow 1962): incumbent firms have less incentive to innovate than new entrants, because incumbents would cannibalize their existing

[1] We thank Mercedes Delagado for sharing patent data on these 25 research-intensive universities.

profits.[2] The AH results highlight the distinction between startups and incumbents within the private sector of the economy: even though both types of firms seek to maximize profit, they pursue meaningfully different kinds of innovation.

The model of creative destruction presented in AH sits in some creative tension with a different "research line" by one of the authors, ADS. The ADS framework models a multi-stage (rather than single stage) research line whereby the optimal organization type for innovation is determined by a wage differential between academia and the private sector. This wage differential is driven by the disutility suffered by the researcher when she is "focused" or forced to perform her non-preferred strategy to advance a research line (Stern, 2004). As a result of this wage differential, high-freedom organizational forms (i.e., academia) are optimal early in the research line, while high-focus organizational forms (i.e., the private sector) take over once the expected value of success from focused effort exceeds the lost utility from the lack of researcher freedom. In this framework, the key insight is that different organizational forms generate different incentives for researchers, and are therefore better-suited to different stages of the innovation process.[3] The ADS model emphasizes the contrast between academia and the private sector, but assumes that all private-sector organizations are identical. It therefore serves as both a counterpoint and a complement to the AH framework, which highlights the differences between startups and incumbent firms but does not consider academia as a source of innovation.

Bringing together AH and ADS, it becomes apparent that a comprehensive understanding of the innovation process and the ecosystems in which innovation takes place must grapple with

[2] As discussed by Schumpeter (1942), it is possible for the tendency toward creative destruction to be overturned if incumbents have access to more effective means of innovation, e.g., because of returns to scale or superior access to funding or scientific talent. We are able to directly assess this possibility in our empirical analysis.
[3] These differences are further accentuated by other aspects of the innovation ecosystem, most prominently the degree of openness: increasing open access to research inputs has a disproportionate positive impact on researchers operating under academic freedom (Murray et al. 2016).

both differences between organizational forms (public and private) as well as heterogeneity among private-sector firms. A natural question arises when combining these theoretical frameworks: when a research line transitions from academia to the private sector (as depicted by ADS), will it be transferred to an incumbent firm or to a startup venture (as depicted by AH)? More broadly, what is the role of startups in moving innovation from ideas to impact, what is their contribution to the development of multi-stage research lines, and when does startup advantage arise?

Motivating Examples: Global Health and Climate Security

Startup ventures are an increasingly important organizational form in the commercialization of university technologies in critical areas such as global health, climate security, and infrastructure. Below we examine three high-profile technologies originating from universities: mRNA vaccines, nuclear fusion, and low-carbon manufacturing. Within each example, we explore the organizational choices made by founders to create startup companies even under conditions where it is not immediately apparent why a new entrant would be the optimal organization to develop the next stages of these innovations. A deep analysis of these cases provides us with key motivating insights into the conditions that drive startup venture innovation, guiding our formal modeling and shaping our empirical analysis of these phenomena.

Global Health: mRNA Vaccines The global health crisis initiated by the COVID-19 pandemic is a clear illustration of the critical role of startups in the innovation process. Both of the early successes in vaccine development were completed by startups – Moderna in the United States and BioNTech (in collaboration with Pfizer[4]) in Germany. While these ventures became household names in late 2020, both were founded approximately a decade previously with the express goal

[4] We emphasize the role of BioNTech rather than Pfizer because BioNTech was heavily involved in the mRNA research that ultimately generated its COVID-19 vaccine, while Pfizer only joined BioNTech's efforts to develop mRNA vaccines in 2018 (focusing specifically on influenza), and its COVID-related contributions focused on providing support with clinical trials, logistics, and manufacturing.

of building on important research lines undertaken in a range of different academic research laboratories around the world (most notably at the University of Pennsylvania, Philadelphia). In each case, the new ventures were continuing research lines that had been long in the making. The first demonstration of mRNA coding for 'therapeutic proteins' in a mammalian setting took place in 1990, and the problem of delivery *in vivo* into cell lines was overcome only in the mid- to late-2000s (and even then, only in lab-based experiments in research mice). The founder-entrepreneurs of BioNTech and Moderna had specific private beliefs (and insights) into the potential value of mRNA vaccines, and were therefore willing to create new ventures and find investors also willing to develop a novel and untested technology. In contrast, incumbent pharmaceutical firms perceived mRNA vaccines to be relatively risky, and chose to remain uninvolved in the early stages of research, only later forming partnerships with the two companies. Specifically, Moderna's first vaccine-related licensing agreement was with Merck in 2015,[5] while BioNTech only became involved with Pfizer through a research collaboration for influenza signed in 2018.[6] Both partnerships focused more on later-stage clinical trials and commercialization, with the startup partner responsible for the initial vaccine development.[7] By highlighting the critical role of startup-developed technologies, COVID-19 has provided a timely window into the essential role of startup ventures in crisis innovation (Johnson and Murray 2021).

Climate Security: Nuclear Fusion Beyond global health and life sciences (where the role of startup ventures is well-known), startups have also made significant contributions to innovation in longer-term challenges such as climate security, based either on unique (private) information

[5] See https://investors.modernatx.com/news/default.aspx for details.
[6] See https://biontech.de/sites/default/files/2019-08/20180816_BioNTech-Signs-Collaboration-Agreement-with-Pfizer.pdf for details.
[7] Interestingly, the cases of Moderna and BioNTech's competing COVID vaccines offer a valuable contrast in commercialization strategies: Moderna chose to pursue commercialization of its mRNA vaccine through an integrated approach, while BioNTech chose to partner with Pfizer for clinical trials, marketing, and distribution.

and beliefs of the founding researchers or on the fact that incumbents underinvest due to their high switching costs and significant inertia. We illustrate these points by contrasting two recent climate security startups – one in energy and the other in manufacturing, both Boston-based and both with funding from The Engine (a venture capital fund built by MIT but now independent, whose emphasis is on investments in startup ventures applying novel and complex science and technology to missions of global significance).

Commonwealth Fusion Systems (hereafter CFS) was founded in 2018 by MIT post-doctoral researcher Robert Mumgaard, then at the MIT Plasma Science and Fusion Center, and offers a critical example of the value of private information in generating a startup advantage. The venture uses only recently (commercially) available high-temperature superconductors to create magnets that enable a new approach to fusion system design. If successful at full scale, the approach will provide the first net-energy-producing fusion reactor ready for widespread deployment. CFS builds upon long research lines that date back to the 1950s: fusion energy research has been funded by governments around the world for decades prior to its founding. Despite this long history, fusion power has always been considered to be "twenty years away" by incumbent energy providers (and even among many academics), suggesting a research line of extensive duration and considerable resource requirements likely only to be taken on by a large industry incumbent with substantial resources. In addition from 2008 onward, U.S. government funding for fusion research was shifting toward a single massive-scale fusion project based in France known as ITER.[8] Nonetheless, by 2015, Mumgaard along with Professor Dennis White and others at MIT believed that the new improvements in superconducting magnets, combined

[8] For more details, see https://www.iter.org/.

with new reactor designs, meant that there was finally an innovation path to commercial deployment of fusion power.

Importantly, the private information and resulting beliefs held by the CFS founders were central to their willingness to pursue this research line from academia into a private sector venture. At the time of founding, it would have been extremely difficult (and costly) for large incumbent energy companies to test these beliefs and to pursue the project internally; in contrast, startup investors were able to share funding risks across a large syndicate (which included Italian energy incumbent ENI) and allocate funding conditional on key technical milestones. As a result, CFS provides a canonical example of settings where private information about the risk of a particular research approach is central to advantages that a startup venture has in the further development of a promising research line, particularly when academic funding is increasingly limited. Moreover, it does so in a context in which incumbent risks of cannibalization are likely very low given the extremely long deployment timelines and the complementarity of zero-carbon electricity generation with existing power transmission infrastructure.

Climate Security: Low-Carbon Manufacturing New ventures also play a role in contexts where cannibalization, rather than private information, is the dominant difference between startup ventures and incumbents. As a case in point, the new venture Sublime Systems is currently commercializing a novel process for low-carbon cement, in an industry where established incumbents have traditionally under-invested in research lines when faced with high switching costs and potential cannibalization. Sublime founder Dr. Leah Ellis developed a new approach to reducing cement kiln emissions by up to 50% (with few if any significant changes to the properties or chemistry of the resulting cement). A significant contributor to climate challenges, cement production today comprises approximately 8% of global CO_2 emissions (Lehne & Preston, 2018).

Nonetheless, cement production has been subject to very little innovation, with its production dominated by a small number of large established firms (Lehne & Preston, 2018) using a traditional process that uses lime, silica, and alumina simply mixed under high temperatures. The innovation under development by Sublime Systems (and initiated at Dalhousie University and then MIT) uses an electrochemical method which electrifies Portland cement manufacturing and in doing so produces concentrated CO_2 that is easily captured.[9] While novel as a research line applied to cement, electrochemistry is an extremely well-established branch of chemistry (founded by, among others, Alessandro Volta at the turn of the nineteenth century) that has been used as the basis for industrial processes since the extraction of aluminum (by Héroult and Hall in 1886) and ammonia (via the Haber process). The lack of incumbent interest in alternative research lines suggests a clear case of incumbent risks of cannibalization rather than any private information about the risks of a new approach. As such, Sublime Systems provides a clear contrast to Commonwealth Fusion: its entry into the cement industry highlights a setting in which startup ventures play an essential role in the innovation process due to their greater willingness to cannibalize existing industry products.

Motivating Trends: AUTM Licensing

Examining trends in university patenting license provides further evidence of the growing startup advantage in commercializing university technologies. According to the data from Association of University Technology Managers (AUTM 2020), illustrated in Figure 1, there is significant growth in the proportion of startup licenses granted by the top-25 research universities in our sample, rising from 17% in 2001 to 29% in 2019. This rise in startup licensing is accompanied by a corresponding decline in the share of licenses to large incumbent firms, which

[9] For the full details of Dr. Ellis's process for the electrification of cement manufacturing, see Ellis et al. (2020).

fell from 37% to 26% during the same period. The proportion of licenses to small, non-startup firms remained relatively unchanged. Importantly, the increasing share of startup licenses comes on top of an 80% rise in the overall rate of university licensing during the same time period, resulting a three-fold increase in the raw number of licenses to startup ventures. As startups displace the role of large firms,[10] it becomes even more important to understand how startups ventures contribute to the larger innovation ecosystem via their unique role in the innovation process.

INSERT FIGURE 1 HERE

Theory and Predictions

The Role of Academic Startups in the Process of Innovation

Startups play a distinctive role in driving innovation during the transition from early-stage research in academia to the private sector. However, as discussed above, there is a gap in the theoretical literature in terms of explaining how and why startups might be able to pursue innovations that are either infeasible or sub-optimal in other institutional contexts. To bridge this gap, we first synthesize the core structure and insights from the models of innovation in AH and ADS, and then turn to the potential role for startups in the context of a theoretical framework that combines the insights of the above models.

Our starting point is the ADS model, which emphasizes how the distinctive allocation of decision rights – most notably, over whether researchers have the ability to choose their own research direction (i.e., "academic freedom") – influences the organization of research between

[10] The trends in university licensing stand in significant contrast to the trend of the rising roles of incumbents in patent activity, as demonstrated in Figure 3 of our own analysis as well as recent work on declining business dynamism (Akcigit & Ates 2021). This contrast is likely due to the differences between university licensing and corporate patenting, where the former tends toward basic research and the development of new technologies, while the latter is more often associated with applied research, incremental innovation, and patent thickets.

academia and the private sector. The ADS model focuses on a multi-stage research process: at each of k stages of a given research line, a decision-maker chooses whether to attempt to move research to the next "stage" of the line (or perhaps explore an alternative direction). If the decision-maker chooses to attempt to advance the research line, a single researcher exerts effort (for which she has to be paid a wage w), and success in moving the project to the next stage occurs with probability p. If all k stages of a research line are executed successfully, the result is a final output with value V; otherwise, the research line returns no value.

The key insight of ADS arises from considering the impact of the allocation of control rights over whether to pursue the given research line; in other words, the identity of the decision-maker is the crucial distinction between the organizational forms of academia and the private sector. On the one hand, if a private-sector manager maintains control, their primary interest will be in projects that have a chance of (eventually) producing commercial value, and so they will always direct the researcher to pursue the next stage of the commercially oriented research line (ADS refers to this as the "practical" research direction). However, if the control over research direction is allocated to the researcher, as it would be in academia, she *might* (with probability α) choose the practical strategy (this research might be of intrinsic interest to her), but with probability $1 - \alpha$, the researcher would choose an alternative research direction to pursue her own interests. In other words, while the manager and researcher have partially overlapping preferences in terms of research direction, the manager must maintain control over the research direction if they want to ensure the highest probability of attempting to move the innovation to the next research stage.

Of course, this structure raises the immediate question of why the manager would *ever* allocate the rights to the research direction to the researcher (since the researcher may divert the project from the practical direction). Building on the empirical evidence in Stern (2004), ADS assume that the researcher values the ability to maintain academic freedom and so would "pay" z in order to maintain control rights over their research direction. As such, if the reservation wage of a researcher who is given academic freedom is R (in other words, this is the wage she would have to be paid to exert effort while also getting to choose her own research direction), then the manager will have to offer at least $R + (1 - \alpha)z$ to hire the researcher to pursue the research line under the manager's control.

The presence of a compensating differential for academic freedom induces a tradeoff in how to allocate the decision rights for the research project. Most notably, for any given research stage i (assuming that it is worthwhile to pursue this research stage regardless of how the decision rights are allocated), the relative cost-effectiveness of allocating the decision rights to the researcher is increasing in z, the preference for academic freedom, and α, the probability that the researcher would choose the practical strategy anyway.

With the above assumptions, the ADS model is able to determine, at every stage of the innovation process, both whether the research line is feasible and what allocation of decision rights is optimal. If we define Π_{i+1} to be the value of successfully completing research stage *i* and moving to stage $i + 1$ (i.e., the research exerts effort in the practical direction and is indeed successful), then as illustrated in Figure 2, the *minimum* value of the project for which academia is viable: $\Pi_{i+1} > \frac{R}{\alpha p}$ and the *minimum* value for which a project is viable in the private sector is $\Pi_{i+1} > \frac{R+(1-\alpha)z}{p}$.

While project viability is an important constraint on the model's parameters, the primary focus of ADS is on the optimal choice between pursuing the research line in academia or the private sector. Figure 2 illustrates this tradeoff as being determined by the combination of the research line's expected value Π_{i+1} and the researcher's disutility z if she lacks control over the research line's direction. Since $z = 0$ implies that the researcher does not place any value on academic freedom, the control rights for research direction will be allocated to the firm. However, as the relative value placed on academic freedom increases, the manager has increased incentive to forego the wage premium $(1 - \alpha)z$ required to maintain control, and instead prefers to cede control of the research direction of the project to the researcher. Of course, the shifting of decision rights to the researcher makes it less likely that the project will be successful in moving to the next stage of the research line: with probability $1 - \alpha$, the researcher will ultimately choose the alternative rather than the practical path. Concretely, this means that the research line will optimally occur in academia if $z > p\Pi_{i+1}$, while the private sector will be the optimal organizational form in all other cases. Thus, for very high values of *z*, the researcher is allocated control (with the corresponding decline in the probability of success of a commercial project).

The central insight of ADS arises from the fact that, for a multi-stage research line, the i+1$^{\text{th}}$ stage of a research line is always more valuable (from the perspective of its prospective financial return) than the i^{th} stage of that research line (i.e., $\Pi_{i+1} > \Pi_i$). Since the commercial value of the research line is only realized by successfully traversing all *k* stages, the value function is increasing when any progress occurs that allows research to be conducted closer to the end of the line. In Figure 2, from the perspective of a given research line, this is equivalent to a shift along the x-axis for a given level of *z*. Crucially, this has the consequence that, for any research project which is feasible (i.e., $\Pi_1 > \frac{R}{p}$), research will either be conducted in the private sector, or if $\Pi_1 > \frac{R}{\alpha p}$ and $z > \frac{R}{\alpha}$, research will first be conducted in an academic environment (freedom for researchers) and then, as long as *z* is not too high relative to the value of the research line, transition at a discrete point towards a more commercial environment (in which the decision rights are allocated to the manager).

INSERT FIGURE 2 HERE

The ADS framework offers an important conceptual contribution and testable theoretical predictions regarding the allocation of decision rights between a more commercially-focused versus more academically-focused research environment. However, the baseline ADS model does not specify the distinct role played by startup firms (relative to incumbents). Specifically, building on the motivating examples from the previous section, the ADS model does not provide a rationale for why "academic" startups (firms founded in and around leading academic research environments) seem to play such a distinctive role in the innovation process. In technology transfer, this phenomenon is termed "the valley of death"- referring to the gap in interest from either industry or academia at intermediate stages of the research process. We argue that overcoming this "valley of death" is the key role of startups in transferring technologies from academia to the private sector.

We therefore extend the ADS model to consider three (non-mutually-exclusive) factors – equity incentives, private information, and cannibalization – that might be associated with a role for academic startups, and then draw out the potential empirical implications of those factors for academic startup innovation (compared to university and established-firm innovation).

Perhaps the simplest potential explanation for the role of academic startups in the process of innovation is that, at some point in the research process, it is important to provide financial incentives for researchers to pursue the practical path. And, indeed, it may be the case that the prospect of significant financial gain might motivate individuals to exit an academic institution or large established firm to found or join a startup company. However, as long as there is ex-ante symmetric information between the (presumably capital-constrained) researchers and private funders, there is no reason that any monetary incentive that would be required to induce effort on the part of the researcher could not be offered within an established company. More subtly, if an equity-oriented research contract depends on achieving success to the next stage of the research line, then the purpose of such incentives should be to encourage researchers to choose the practical direction. Within the ADS framework, such an incentive contract must be preferred by both the manager and researcher, relative to simply allocating decision rights to the manager. However, as long as the researcher is even slightly more risk-averse than the manager (which seems like a reasonable assumption), then the researcher would prefer to be guaranteed an upfront wage payment (including a premium for ceding control over research direction), and the manager would prefer to guarantee that the researcher pursues the practical strategy. In other words, though academic startup founders do receive equity incentives as part of the founding process, equity incentives on their own do not provide a sufficient rationale for academic startups as a distinctive institutional form in the process of innovation.[11]

A more promising potential driver of academic startups may lie in the fact that the research process itself not only requires effort and choice but also yields private information to researchers. Specifically, for a research project conducted within the academic sector, the researchers themselves are able to observe private signals about the potential value of the overall project or the likelihood of success. For example, in a case such as Commonwealth Fusion Systems, the researchers are uniquely informed about both the nature of the technical breakthrough they have developed within the university, and also the challenges they are likely to confront in the overall process of innovation. From the perspective of the ADS model, the researchers themselves might receive a private signal of the value of p for future stages of the project, and so know whether continued effort is likely to ultimately yield a successful innovation. For example, one could

[11] For an in-depth discussion of the incentive structures most conducive to motivating innovation, see Manso (2011).

consider an extreme case where researchers were able to observe whether $p = 1$ for all future stages of the project (in other words, the project is guaranteed to succeed if subsequent research is directed towards the practical path). Of course, the researcher could disclose their information about that signal to the manager; however, under the baseline ADS model, the researcher's wage following disclosure would simply be equal to $R + (1 - \alpha)z$, as the manager would capture the surplus of the successful project. By contrast, if the researcher were compensated using equity, or chose to found a separate organization to commercialize the innovation, she could capture the some or all of the value of the research line's output, V, net of any wage costs for commercialization in future periods. More formally, in cases where researchers receive private information regarding the probability of success p of their research line, and this probability is sufficiently high relative to ex-ante expectations, then an organizational form that offers the researcher either partial or full ownership of the research line's payoffs through equity-based compensation would dominate both academic institutions and established firms.

Shifting from theory to theoretical predictions, would anticipate that academic startups are more likely in the case where the academic researcher receives private information about the underlying value or potential of success of the research line. This ability of researchers to select the most promising projects on the basis of private information implies that high-risk, high-reward research lines will preferentially flow into startups rather than incumbent firms. This same selection effect, in combination with the tendency of research lines to increase in value as they advance from stage to stage, also implies that startup innovation will also be more novel and more impactful than university innovation. We formalize these insights in the following hypothesis:

Hypothesis (H1). *Vision: Relative to both universities and established firms, startups will be associated with innovation that is more original and more impactful over time*

Our third potential driver of startup advantage is based on the foundations of creative destruction, as presented by AH. Specifically, we would expect there to be differences in the perceived value of a research line when comparing the perspective of a startup to that of an established firm. An established firm that brings a new product or service to market would tend to cannibalize at least some of its existing offerings, so its marginal benefit of successfully completing the research line would be lower in proportion to the degree of cannibalization. By contrast, startups have no pre-existing products or services to cannibalize, and would therefore obtain a greater marginal benefit from completing the same research line. Within our theoretical framework, we formalize this idea as the case where $V^{startup} > V^{incumbent}$. Importantly, because of the recursive nature of multi-stage research lines, this implies that $\Pi_{i+1}^{startup} > \Pi_{i+1}^{incumbent}$ for all stages *i* of the research line. The degree of cannibalization leading to this discrepancy in perceived value would be greater for highly-novel and disruptive technologies, since they would likely by substitutes for existing products and services offered by incumbents, as in the example of the cement-production technologies developed by Sublime Systems.[12] By contrast, the case where innovations would tend to complement incumbents' existing offerings, is captured by having $V^{incumbent} > V^{startup}$. Indeed, the case where the research line's output is more valuable to incumbents than to startups is discussed at length by Gilbert and Newbery (1982), who focus on the decline in industry profits if a monopolist allows a new entrant to join the market. Importantly, this would apply only to innovations that would share the market with existing products (i.e., incremental technologies), rather than innovations that would replace the incumbent's offerings (i.e., radical technologies) as described in Arrow (1962) and AH. This distinction between incremental and radical technologies is further explored by Gans and Stern (2003), who highlight the importance

[12] We would also expect that academic researchers may care about the non-monetary impact of their inventions, particularly by placing more weight on the broader social value of a potential invention.

of complementary assets and functional "markets for ideas" as key drivers of whether startups are likely to cooperate with incumbents or seek to undermine them as they develop new technologies.

In both of the above cases, any difference in the perceived value of a completed research line would flow back to earlier research stages, based on the private-sector institution developing the research line. By introducing AH's creative destruction into our framework, we generate a second hypothesis in the context of academic innovation ecosystems: we would expect to see more radical and disruptive innovations pursued by startups, while more incremental innovations are pursued by established firms.

Hypothesis (H2). *Creative Destruction: Relative to established firms, startups will be associated with innovation that is more radical and more impactful over time*

Setting, Methods, and Data

The goal of our empirical analysis is to test the above hypotheses by documenting the unique characteristics of startup patents, and to contrast them with patenting activities performed within academic research institutions and private-sector organizations. In this section, we describe our empirical setting of top-research-university ecosystems, define key variables, and characterize our empirical strategy.

Empirical Setting

Our empirical setting targets the regional ecosystems around the 25 top research universities in the United States, as defined by Delgado and Murray (2021). Specifically, these are the universities that were listed as assignees on the largest number of patents granted in the 2011-2015 time period. While this sample is selected partially for convenience, it is worth noting that these universities make up approximately 50% of all academic patenting during our sample period. In addition, they account for approximately 40% of all university patent licensing as captured by

AUTM records, reflecting the strong emphasis on both innovation and commercialization for the universities in our chosen sample.

In addition to the geographic focus above, we also target the time period of 2000-2015 for our analysis, again coinciding with the sample period of Delgado and Murray (2021). This time period also overlaps with the decline of large firms and the rise of small firms and startups in university licensing, as described in our analysis of AUTM licensing data above. The termination of the sample period in 2015 is necessary in order to have sufficient time to collect forward citations following patent publication. Moreover, the 2000-2015 time period also overlaps with the longer-term expansion in patenting (Lerner and Seru, 2021), and the increasing importance of agglomeration economies that tend to be co-located with major research universities (Carlino and Kerr, 2015). In light of these trends, our empirical setting offers insight into a narrow but crucially important aspect of the innovation ecosystem.

Data Sources and Key Variables

The primary data source for our analysis is the PatentsView database, which tracks all patents granted by the U.S. Patent and Trademark Office from 1976 to the present. This database includes a wide range of patent-level measures, including filing and grant dates, primary technology classes based on the Cooperative Patent Classification (CPC) system, and forward and backward patent-to-patent citations. PatentsView also disambiguates patent assignees and their locations, as well as the locations of inventors listed on each patent. We supplement this dataset in a number of important ways: first, using the dataset constructed by Delgado and Murray (2021), we identify the patents granted to the 25 top U.S. universities by patent count, based on the 2011-2015 time period. Second, we merge in non-public administrative data from the USPTO to identify patents processed under small-entity or micro-entity status, as reflected by patent fees. Finally, using the precise latitude and longitude of each of the universities in our sample, we are able to identify the full population of granted patents that have at least one assignee and at least one inventor within a ten-mile radius.[13]

Using the data described above, we proceed to construct the central variables of our analysis. Our primary explanatory variable is *OrganizationType*, a patent-level categorical variable which takes the values of "University," "Startup," or "Incumbent." University patents are directly identified by Delgado and Murray (2021). Startup patents are defined as those processed under small-entity or micro-entity status (indicating that they have no more than 500 employees), and patents where all assignees are at most five years old in terms of their "patent age," or years from their first patent filing.[14] Finally, incumbent patents are those that were processed under normal (i.e., undiscounted) patent fees, and where all assignees are at least ten years old in terms of patent age. Any patents that do not meet these definitions are excluded from our analysis, in order to highlight the differences between these groups. For example, a patent that results from a collaboration between a startup and an incumbent would not be classified under either category. In extended analysis, we also introduce an additional category for *OrganizationType*: the "Scale-Up," which we define as a former startup that continues to generate patents but has either stopped qualifying for small-entity or micro-entity status, or is now older than five years in "patent age." We use this measure to capture the degree of creative destruction within our sample, as we track the displacement of incumbents by both startups and scale-ups.

[13] For our regression analysis, we also merge in county-year covariates including population, income per capita, unemployment rate, labor force participation, and the total number of patents granted within the county for each calendar year.

[14] We calculate "patent age" based on the gap between filing dates for each assignee's patents. However, because our outcome measures are often based on citations, we group patents by grant year in our figures and regression analysis.

Our outcome variables focus on capturing the differences in the type of patents being produced by the organizations described above. Our first set of outcomes track the number of forward citations (i.e., citations from future patents) that accrue after a given patent is published. In order to capture the dynamic aspects of forward citations, and to effectively compare patents of different vintages, we divide forward citations into three distinct time periods relative to the original grant date: 0-5 years, 6-10 years, and 11-15 years. The last of these categories, in combination with the fact that we track patent citations through the end of 2020, means that we are limited to patents granted no later than 2005 when we analyze the full range of forward citation dynamics. In addition to average citation counts, we also track outlier patents. For this outcome measure, we calculate the 95^{th} percentile of forward citations (at both five- and ten-year horizons) for each combination of grant year and CPC class (e.g., grant year 2004 and CPC class A01). We then generate indicators for whether each patent exceeds these percentiles, thereby falling in the top 5% of all patents for its year and technology.

Moving beyond citation rates, our final set of outcome measures seeks to capture the nature of innovation by different organizational types by looking at the distribution of citations across patent classes. To do so, we replicate the patent-level measures of *originality* and *generality* first developed by Hall, Jaffe, and Trajtenberg (2001). The *originality* measure tracks the dispersion of backward patent-to-patent citations across technology classes, and aims to capture more novel and groundbreaking patents. Similarly, *generality* tracks the dispersion of forward patent-to-patent citations across technology classes, and aims to capture how basic or fundamental a patent is. For this latter measure, we once again limit the time horizon of forward citations to facilitate comparisons across patents of different vintages, calculating generality on the basis of both five-year and ten-year forward citations.

Methods

Using the variables described above, our methodology is simple and straightforward: we calculate averages and standard errors for our outcome variables across organizational types, and plot the results across a series of bar graphs. This approach does not offer causal identification: we are not able to estimate a counterfactual for how a given innovation would have differed if it had been developed under a different organizational form. However, the strength of this approach is that it highlights the stark differences in patented innovations across universities, startups, and incumbents in terms of their respective contributions to their regional innovation ecosystems. In extended analysis presented in the appendix, we replicate the comparisons in our figures using a multiple-regression framework that controls for a broad range of covariates including fixed effects for county, grant year, and technology, and a range of additional covariates. While the regression-based methods are also only able to capture correlation rather than causation, the breadth of control variables allows us to rule out alternative explanations like startups and incumbents working in different technological fields or operating in different geographic locations. Overall, our methods aim to highlight the differences among patents emerging from universities, startups, and incumbents, while focusing on locations and time periods where all three groups would be important contributors to the innovation ecosystem.

Results

Our main results begin with Figure 2, which depicts the relative numbers of patents granted to each type of organization in our sample during the 2000-2005 time period. We observe that while universities and startups both generate similar numbers of patents, the majority of inventions in our sample are generated by incumbent firms. This is not a surprising pattern: our sample includes major innovation ecosystems such as Silicon Valley, Seattle, and Boston, where many large corporations have extensive research and development facilities. Having established this baseline, we proceed to examine whether there is a difference in the type of inventions being produced by the organizations in our sample.

INSERT FIGURE 3 HERE

Forward Citations

The primary empirical test of our analysis is to compare the value of the average invention generated within each of our organizational forms. We do so through the measure of average forward citations, which we depict in Figure 3. By restricting our analysis to the 2000-2005 cohort, we are able to track forward citations at horizons of 0-5, 6-10, and 11-15 years from the grant date of the original patent. We find that across all time periods, startup patents generate more forward citations on average than those of either universities or incumbent firms. Even more importantly, this "startup advantage" increases over time: startup patents receive 20% more citations from 0-5 years, 50% more citations from 6-10 years, and 100% more citations from 11-15 years. We replicate these patterns using a multiple-regression framework in Table A1 of the Appendix, which includes a broad range of fixed effects and covariates. Our findings indicate that startup inventions are more valuable and more impactful than inventions emerging form other organizational forms, and that this advantage only increases over longer forward-citation horizons.

INSERT FIGURE 4 HERE

In addition to the average value of patents produced by different organizations, we are also interested in the rate of outlier inventions. For this analysis, we turn to Figure 4, which presents the rates of patents in the top 5% of their grant-year and technology cohort based on forward citations. It is worth noting that all of the rates in Figure 4 exceed 5%, reflecting the fact that inventions generated within the ecosystems of top research universities tend to be more valuable than the broader patent population. Turning to a comparison across organizational types, we find that startup patents are approximately 40% more likely to be "outlier inventions" in the top 5% of the citation distribution. This is true whether we look at five-year or ten-year forward citation horizons. Further, startup patents are more likely to be outliers when compared not only against incumbents, but also against university inventions. This latter finding is particularly surprising, as the top research universities in our sample are generally viewed as generating highly-influential and valuable research. In extended analysis depicted in Figure A1 of the Appendix, we also show that the best patents generated by startups are far more impactful than the best patents generated by either universities or incumbents, and that this advantage is even more pronounced at longer forward-citation horizons. Overall, these results offer strong evidence for the existence of a "startup advantage" in terms of the value of the inventions they generate: not only are startup patents more cited on average, but they also are more likely to be outliers that have an outsize contribution to future innovation, and over longer horizons than inventions from either incumbents or universities.

INSERT FIGURE 5 HERE

Originality and Generality

Beyond exploring the value of inventions through the analysis of forward citation rates, we would also expect to find differences in the type of innovations generated by startups. For this analysis, we explore the measures of originality and generality, which are calculated based on the dispersion of backward and forward citations, respectively, across patent classes. Figure 5 presents our findings: for both originality and generality, we find that startup and university inventions are nearly identical, and that both exceed the invention s of incumbents in these measures. As these measures are both based on Herfindahl calculations, they always range between zero and one; in our sample, both measures also have a standard deviation of approximately 0.28. In light of this, startup inventions are more original by approximately half a standard deviation relative to incumbent innovations, and more general by approximately one-third of a standard deviation for both five- and ten-year forward citation horizons. We find a similar pattern of results with slightly smaller differences when implementing multiple-regression estimators in Table A2 of the Appendix. Overall, our results in this section suggest that startup innovation is quite similar to the inventions generated by university researchers, and that both of these organizations generate patents with greater originality and greater generality than those of incumbents.

INSERT FIGURE 6 HERE

Assignee Dynamics and "Scale-Ups"

Having highlighted a number of significant differences between startup inventions and those of universities and incumbents in the previous sections, we now turn to the question of dynamics: what happens as the startups we track increase in size and age to become more like the incumbent firms against which they are compared? For this analysis, we introduce the category of "Scale-Up" organizations, which are former startups that are now older than five years in "patent age," or whose patents no longer qualify for small-entity patent fees. We present our main findings in Figure 6, which tracks this expanded set of four organizational types in terms of their citation-weighted patent output over time. We find that scale-ups, which are ruled out by definition in the first year of our sample, grow in importance over time and end up overtaking incumbents in terms of citation-weighted patents by the end of our sample period in 2015. Other organizational types also experience meaningful growth over this same time period, reflecting the overall increase in innovation occurring in these ecosystems. However, the greater level of growth demonstrated by scale-ups indicates that a significant amount of creative destruction is being driven by the former startups identified in our sample. Interestingly, in extended analysis presented in Figure A2 of the appendix, we find that forward citations for these organizational types are approximately equal, and are significantly higher than those for either universities or incumbents. This suggests that the "startup advantage" we identified earlier carries over as startups become scale-ups, and raises the possibility that this higher level of patent value may be an important driver of creative destruction.

INSERT FIGURE 7 HERE

Discussion and Conclusions

In this paper, we present theory and evidence for the unique role of startups in the innovation ecosystem. We begin with vignettes of prominent startups: Moderna, BioNTech, Commonwealth Fusion, and Sublime Systems demonstrate the tendency for startups to pursue novel, high-upside technologies, while Sublime Systems offers an example of a startup aiming to displace existing technologies in an industry dominated by powerful incumbents. We also present evidence from university licensing that startups have taken an increasing share of the technology transfer ecosystem since 2000, reaching nearly 30% of all licenses in 2019 while displacing large firms. Next, combining and extending the insights of AH and ADS, we argue that there are two specific channels through which theory predicts startups will have an advantage in multi-stage research lines. First, startups will have private information on the value of their ideas, and this asymmetry is likely to be highest in novel and high-upside technologies, leading them to be developed internally rather than transferred to an outside party. Second, startups do not face the threat of cannibalization, and are more likely to pursue disruptive technologies that would undermine the competitive advantages of existing incumbents.

We evaluate the above predictions using a dataset of all patents granted within the regional innovation ecosystems centered on the top 25 US research universities from 2000 to 2015. Our results suggest that there is a clear startup advantage in terms of both the average and the right tail of innovation value, as measured by forward citations, relative to both academic and incumbent patents. We then show that startup inventions are both more original and more general than those of other institutions, and conclude by highlighting how former startups, or "scale-ups," have overtaken pre-existing incumbents in their citation-weighted inventions in the space of just fifteen years. Taken together, our findings highlight that startups produce inventions that are both more valuable and more disruptive and novel than those pursued by either private-sector incumbents or academic institutions.

Limitations and Generalizability

While our findings shed light on the significant differences between startup and incumbent innovation, there are several limitations worth noting when considering their appropriate interpretation. First and foremost, our analysis does not offer identification of causal effects: we can only document the trends and patterns that exist within the innovation ecosystems covered by our sample. Thus, rather than attempting to estimate what would have happened if a startup technology were counterfactually developed by an incumbent, we instead show that fundamentally different types of inventions are being pursued within these distinct organizational types. We do not attempt to distinguish between interpretations based on the causal impact of different organizational types from those based on selection effects applied to research lines, innovators, or funding sources. Indeed, we expect these to differ significantly across our organizational types, and seek to document the combined impact of all of these factors. For example, our findings of a "startup advantage" are likely to be driven in part by a selection effect: whether due to private information, credit constraints, or other factors, startups tend to pursue a smaller number of high-impact innovations, while incumbents develop a larger number of innovations that are, on average, lower in terms of forward citations, originality, and generality. While we do not attempt to separate the impact of such selection from the potential for startups to positively impact the technologies they pursue, we do show evidence in Figure A1 of the appendix that the "best" (i.e., 95^{th}-percentile) startup inventions are more highly-cited than the "best" incumbent innovations. This difference implies that at least some of the startup advantage that we document is a fundamental difference in the innovative process, rather than startups selecting only innovations in the right tail of a similar distribution as the one available to incumbents.

Despite these limitations, our regression analysis does include fixed effects for locations, calendar years, and technology classes, and sub-sample analyses indicate that our results are likely to generalize across a broad range of geographies, time periods, and economic sectors. Further, while our sample focuses on university ecosystems with exceptionally high rates of innovation, our broader examination of AUTM data suggests that other ecosystems are likely to experience similar patterns. Overall, we expect the startup advantages in citations, originality, and generality to hold in any setting where startups represent a meaningful portion of the innovation ecosystem.

Policy Implications

By clarifying the unique contribution of startups to their innovation ecosystems, our work suggests a number of policies that would be appropriate when seeking to support and enhance innovation. A prominent issue in innovation policy has been the challenge of increasingly-long gaps between the initiation of a research line and its ultimate commercialization, which can lead to significant under-investment relative to shorter-term research (Budish, Roin, & Williams 2015). How can policy-makers ensure that the maximal number of socially valuable lines are kept alive? Our findings suggest that long-horizon funding for basic research can be a critical driver of startup creation, which in turn generates the startup advantages described in our empirical analysis.

Specifically, university technology transfer could amplify the benefits of startups' superior innovative performance by placing greater emphasis on the potential for long-term value creation and optionality rather than short-term revenue goals.

Breaking down our two channels of startup advantage, we first turn to the concept of "vision" and the ability of startups to pursue high-risk, high-payoff outcomes where they have private information. Given that most research lines depend at least in part on the unique insights of researchers, what type of institutions and policies encourage pursuit of individual vision? We expect that institutions that provide training and support for not only research but also commercialization and scaling would be most likely to generate the benefits of the "vision" channel. Further, this support would complement sources of patient and independent capital to help provide critical "tests" for unique contrarian visions.

Our second channel focuses on creative destruction: Given that the pursuit of research lines from initial idea to ultimate commercialization depends on market incentives, how do we ensure the pursuit of socially valuable innovations that may disrupt existing incumbents? Again, our analysis suggests a number of potentially-valuable policy interventions. First, in light of the prominence of startups in university licensing and technology transfer, interventions targeted at universities, with an emphasis on educating innovators on the path from "idea to impact" (Frølund, Murray, & Riedel 2018). Second, policy-makers can seek to offer signals (and incentives) for research directions that are valuable even in the face of powerful incumbents. Indeed, the incumbents themselves might also be spurred to greater levels of innovation by the rise in competition (Aghion et al. 2005). Closely related to this is the idea of a proactive and pro-innovation competition policy: reducing entry barriers and limiting the power of incumbents to isolate themselves from competition would increase the tendency for innovative entry and subsequent startup advantage. Taken together, these policy interventions are likely to promote higher rates of startup formation and invention, and would either take advantage of or actively increase the startup advantages that we document, and would offer a significant enhancement of the broader innovation ecosystem.

References

Akcigit, U., & Ates, S. T. (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, *13*(1), 257-98.

Arrow, K. (1962). Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors* (pp. 609-626). Princeton University Press.

Aghion P, Bloom N, Blundell R, Griffith R, Howitt P (2005) Competition and Innovation: an Inverted-U Relationship. *The Quarterly Journal of Economics* 120(2):701–728.

Aghion, P., Dewatripont, M., & Stein, J. C. (2008). Academic freedom, private-sector focus, and the process of innovation. *The RAND Journal of Economics*, *39*(3), 617-635.

Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, *60*, 323-351.

AUTM US Licensing Activity Survey: 2020, editors Grant Allard, John Miner, Dustin Ritter, Paul Stark and Ashley Stevens.

Budish, E., Roin, B. N., & Williams, H. (2015). Do firms underinvest in long-term research? Evidence from cancer clinical trials. *American Economic Review*, *105*(7), 2044-85.

Carlino, G., & Kerr, W. R. (2015). Agglomeration and innovation. *Handbook of regional and urban economics*, *5*, 349-404.

Delgado, M., & Murray, F. (2021). *Mapping the Regions, Organizations and Individuals that drive Inclusion in the Innovation Economy* (No. c14571). National Bureau of Economic Research.

Ellis, L. D., Badel, A. F., Chiang, M. L., Park, R. J. Y., & Chiang, Y. M. (2020). Toward electrochemical synthesis of cement—An electrolyzer-based process for decarbonating CaCO3 while producing useful gas streams. *Proceedings of the National Academy of Sciences*, *117*(23), 12584-12591.

Frølund, L., Murray, F., & Riedel, M. (2018). Developing successful strategic partnerships with universities. *MIT Sloan management review*, *59*(2), 71-79.

Gans, J. S., & Stern, S. (2003). The product market and the market for "ideas": commercialization strategies for technology entrepreneurs. *Research policy*, *32*(2), 333-350.

Gilbert, R. J., & Newbery, D. M. (1982). Preemptive patenting and the persistence of monopoly. *The American Economic Review*, 514-526.

Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2001). The NBER patent citation data file: Lessons, insights and methodological tools.

Johnson, E., & Murray, F. (2021). What a Crisis Teaches Us About Innovation. *MIT Sloan Management Review*, *62*(2), 58-65.

Lehne, J., & Preston, F. (2018). Making concrete change: Innovation in low-carbon cement and concrete. *Chatham House Report, Energy Environment and Resources Department: London, UK*, 1-66.

Lerner, J., & Seru, A. (2021). The Use and Misuse of Patent Data: Issues for Finance and Beyond. *The Review of Financial Studies*.

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

Murray, F., Aghion, P., Dewatripont, M., Kolev, J., & Stern, S. (2016). Of mice and academics: Examining the effect of openness on innovation. *American Economic Journal: Economic Policy*, *8*(1), 212-52.

Schumpeter J.A. (1942) *Capitalism, Socialism and Democracy* (McGraw-Hill, New York).

Stern, S. (2004). Do scientists pay to be scientists? *Management science*, *50*(6), 835-853.

Figure 1: Proportion of University Licenses by Firm Type

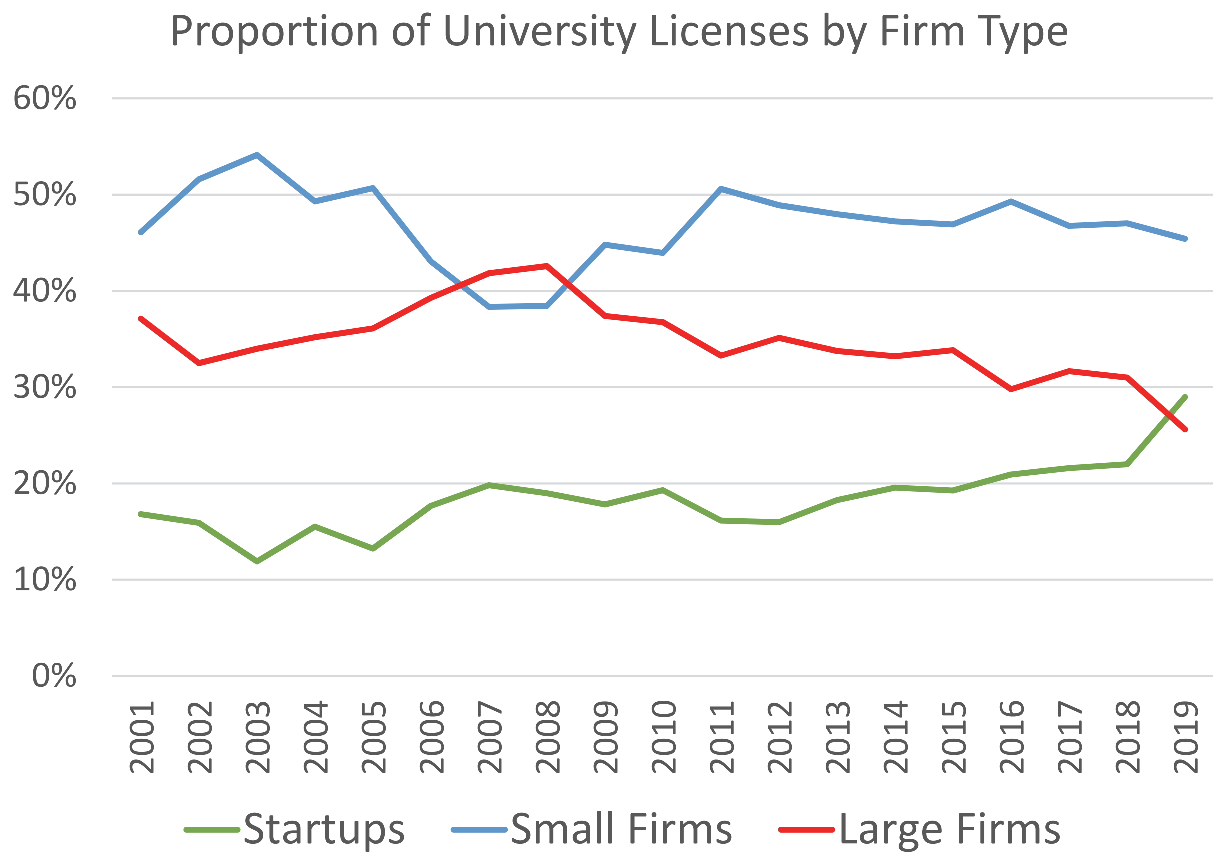

Figure 2: Research Lines in the Aghion-Dewatripont-Stein (ADS) Framework

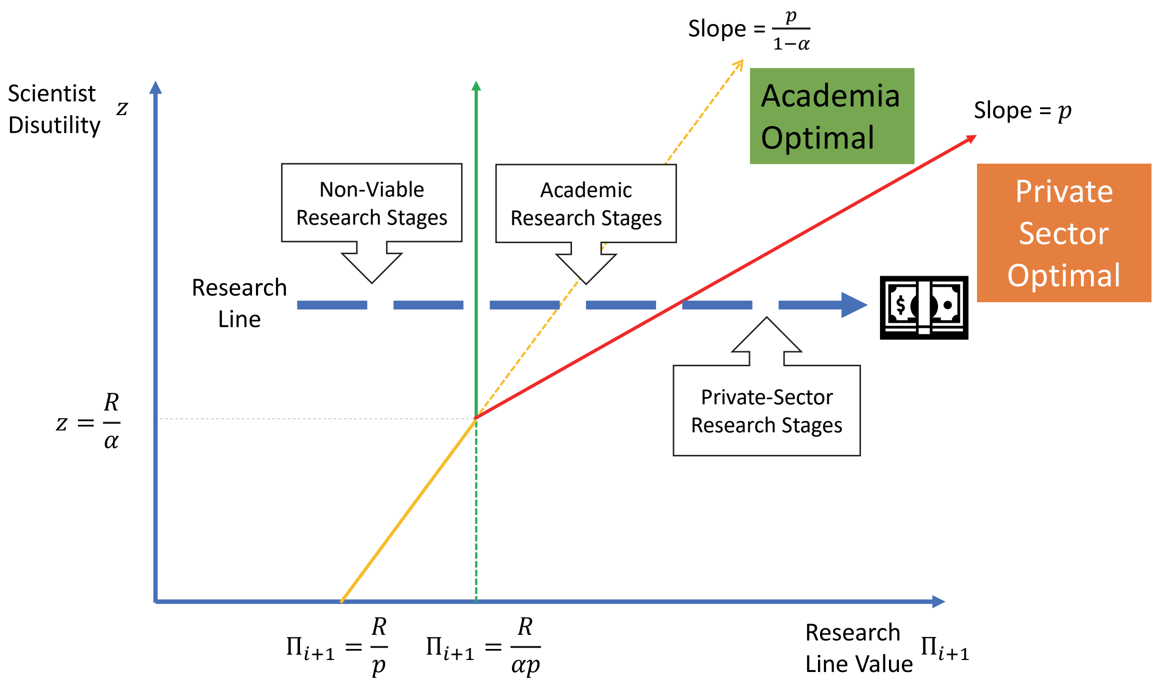

Figure 3: Total Granted Patents (2000-2005)

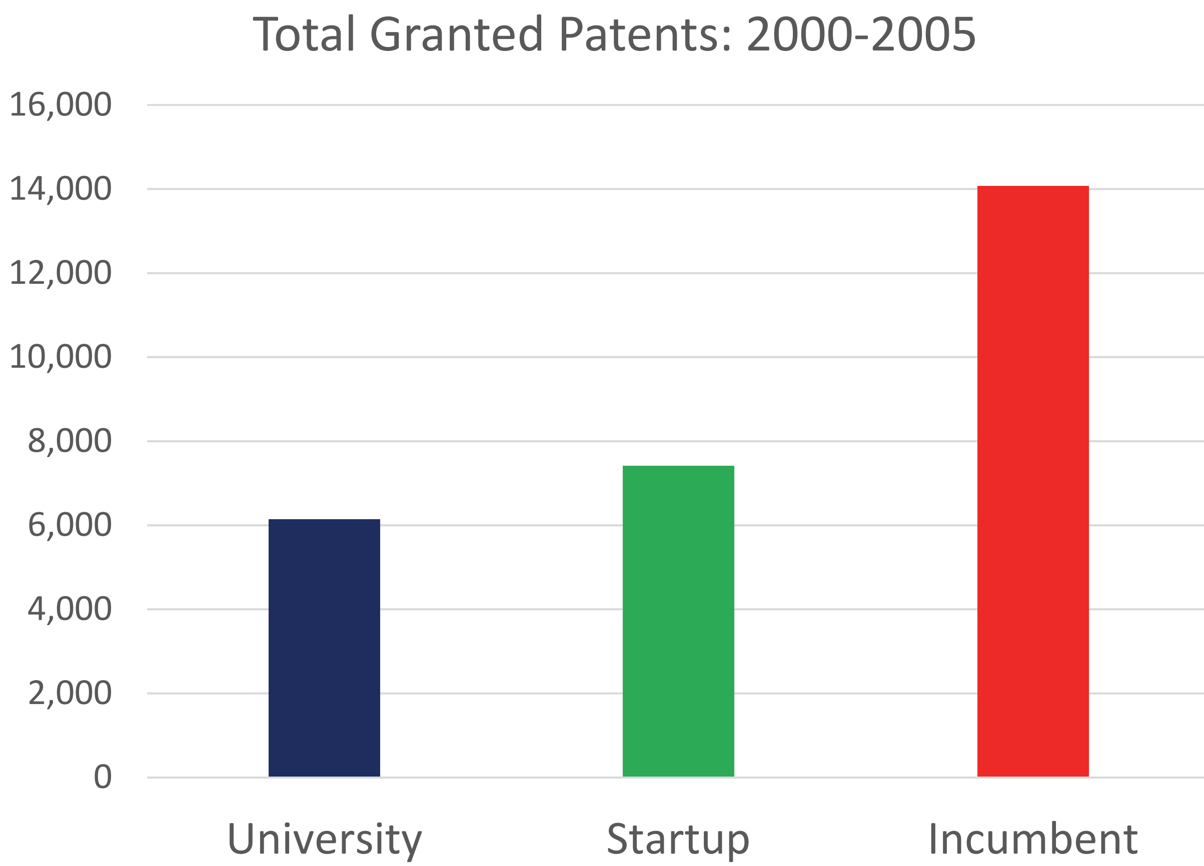

Figure 4: Startup Advantage: Forward Citations

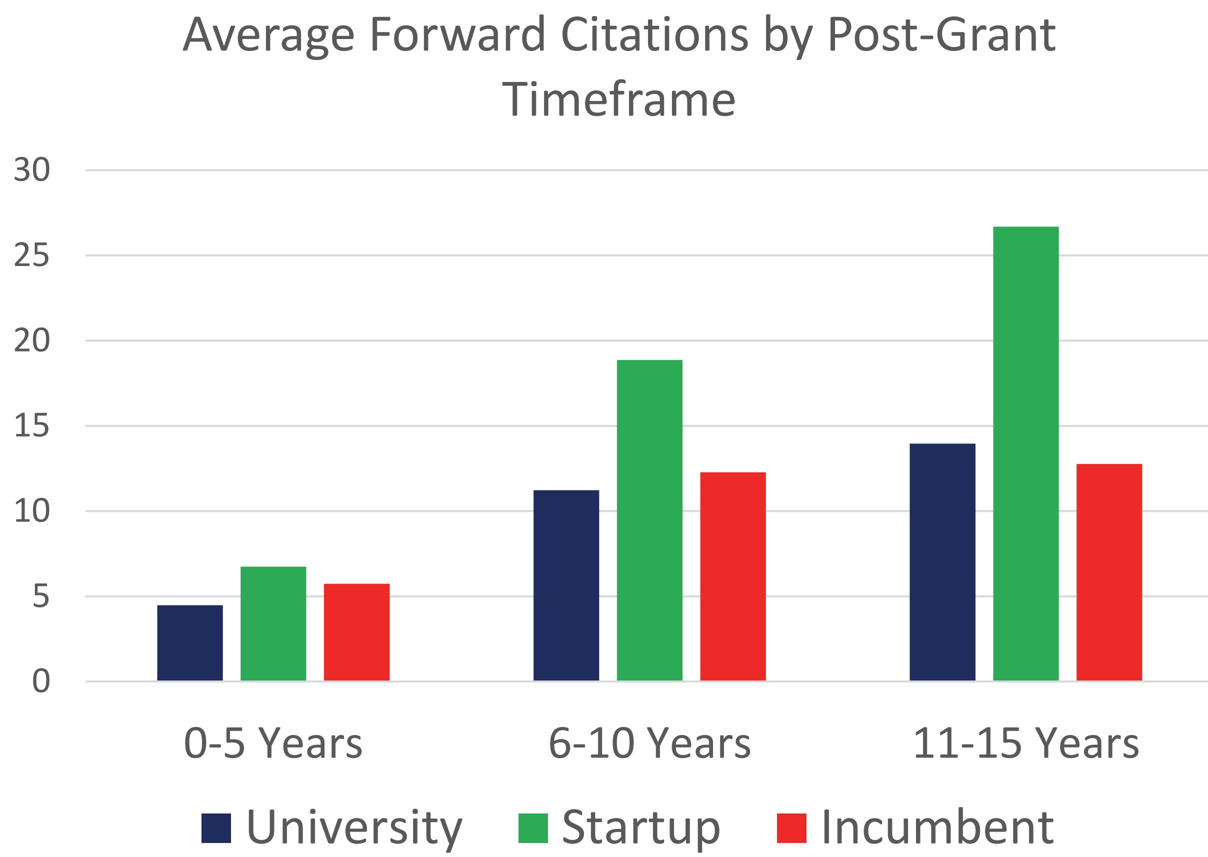

Figure 5: Outlier Patents

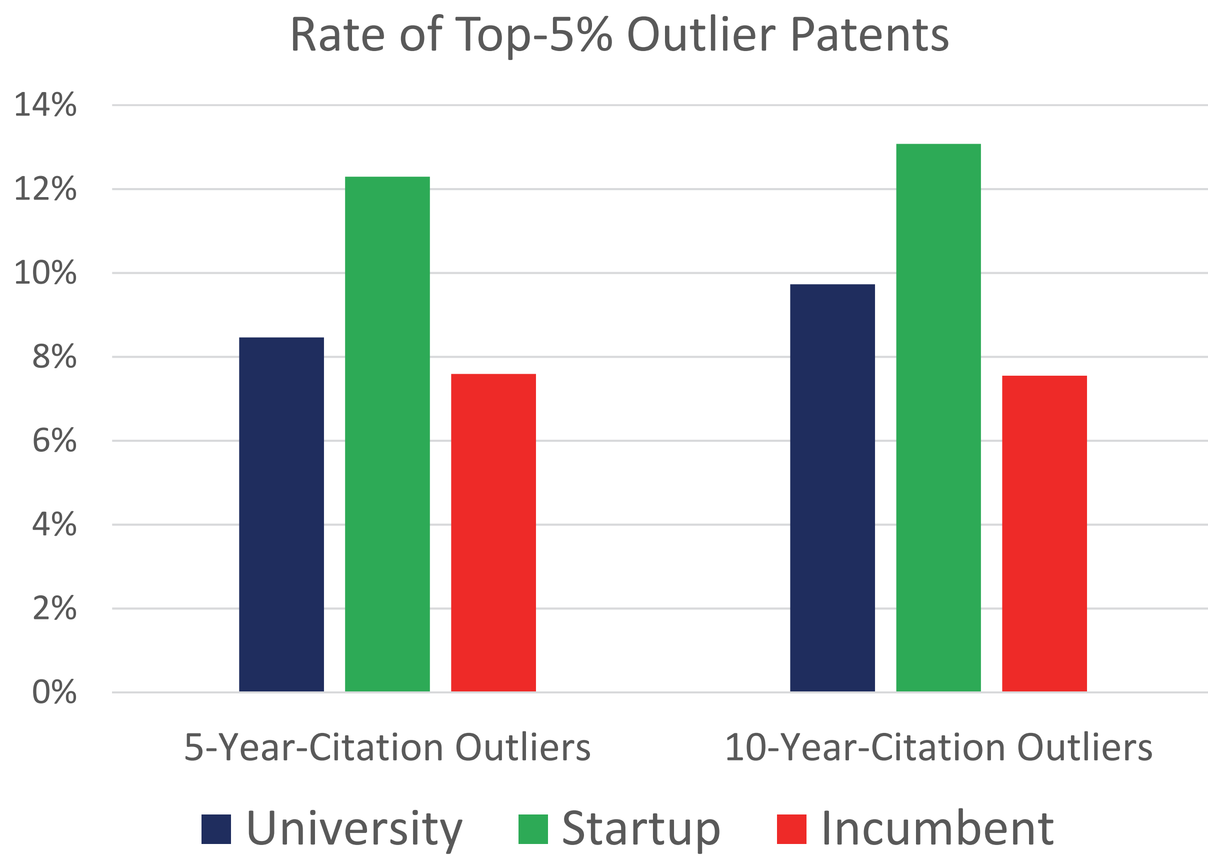

Figure 6: Patent Originality and Patent Generality

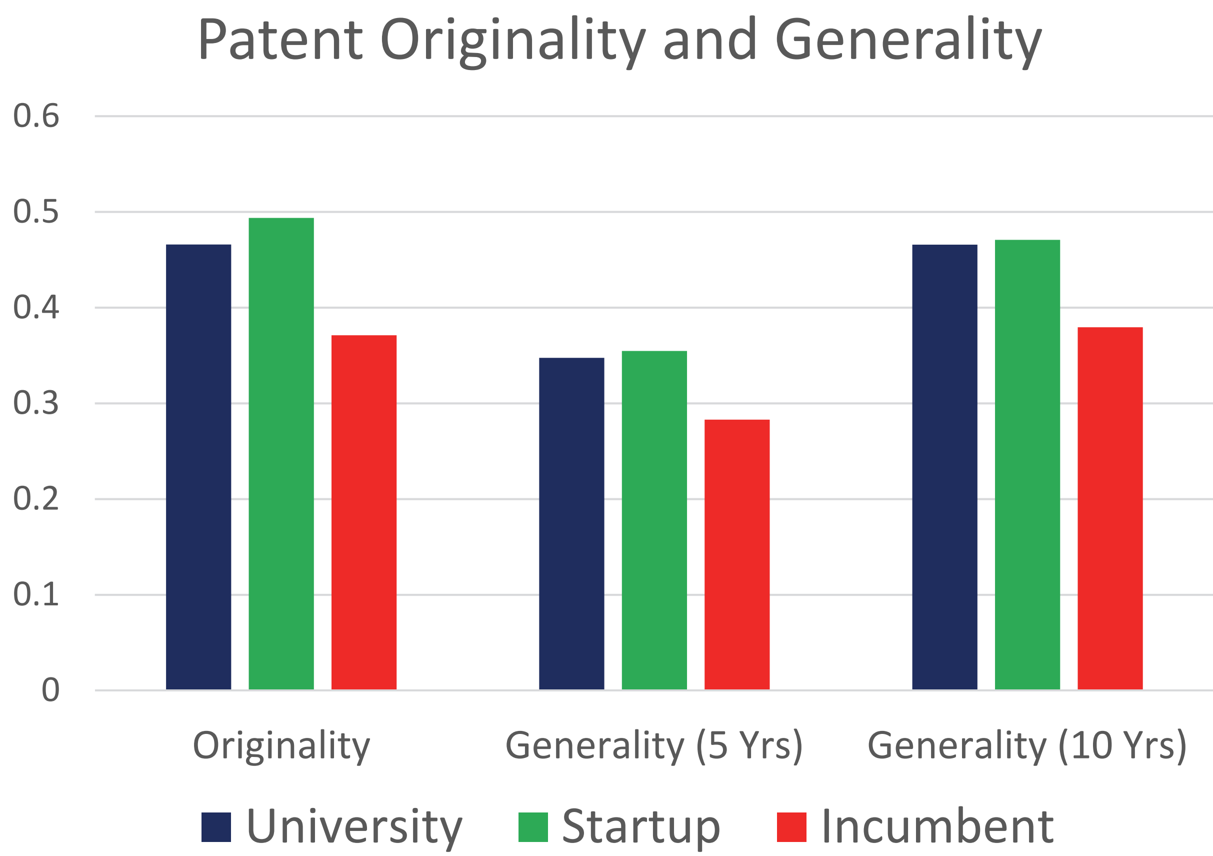

Figure 7: Patent Count Dynamics with Scale-Ups

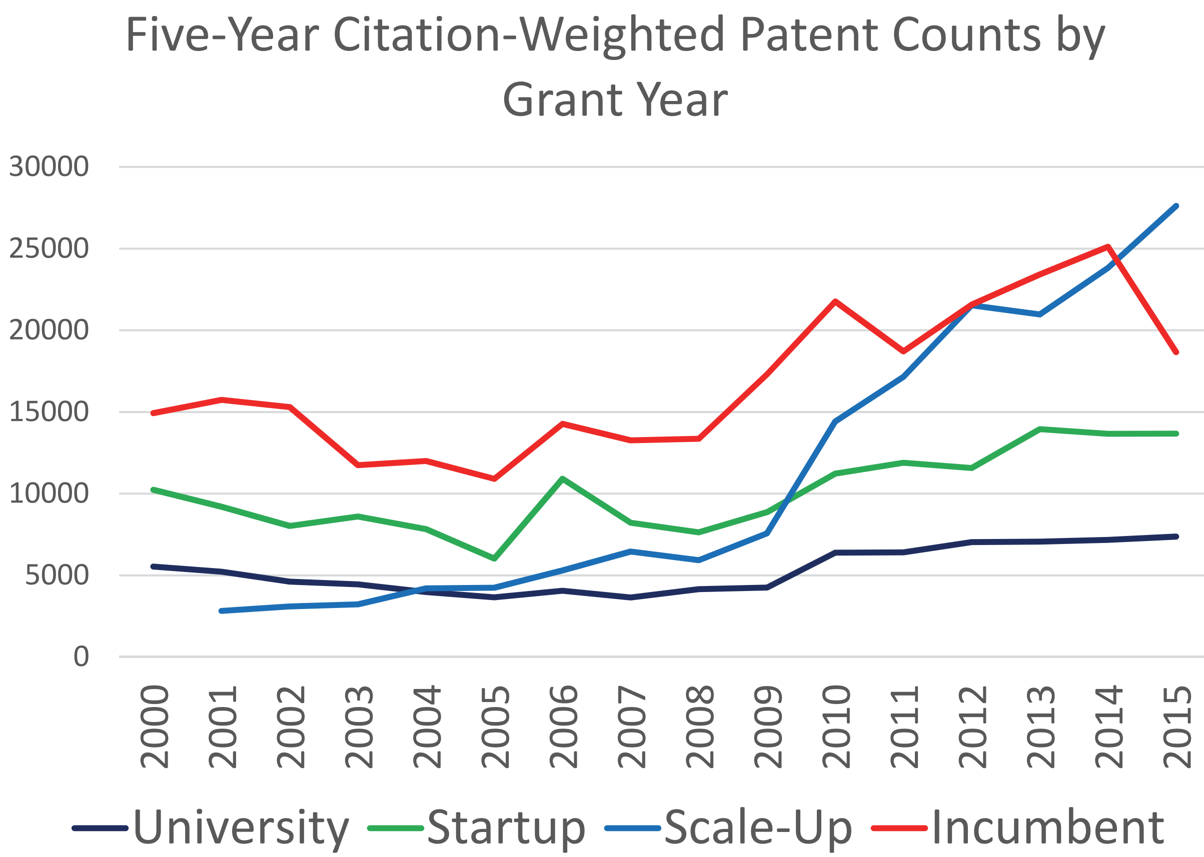

Appendix

Figure A1: 95th Percentile of the Forward Citation Distribution

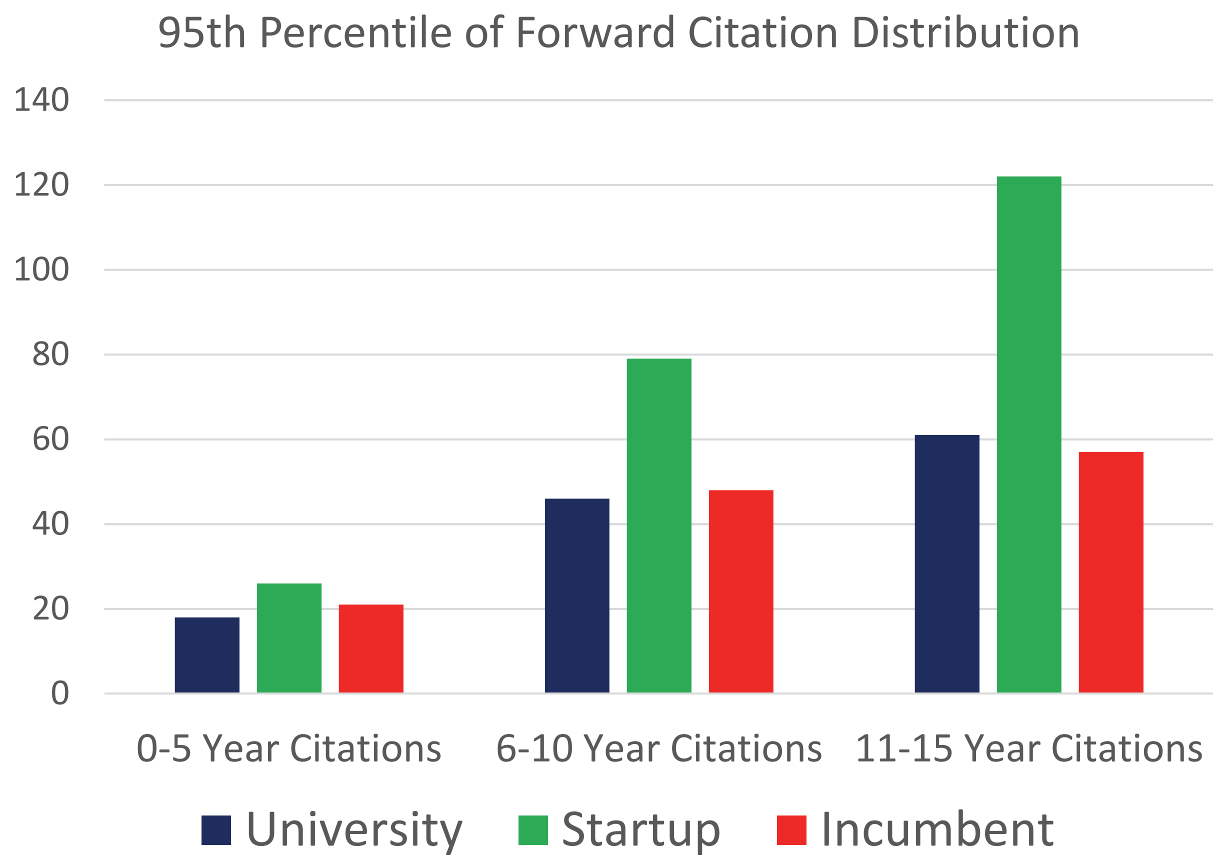

Figure A2: Forward Citations with Scale-Up Patents

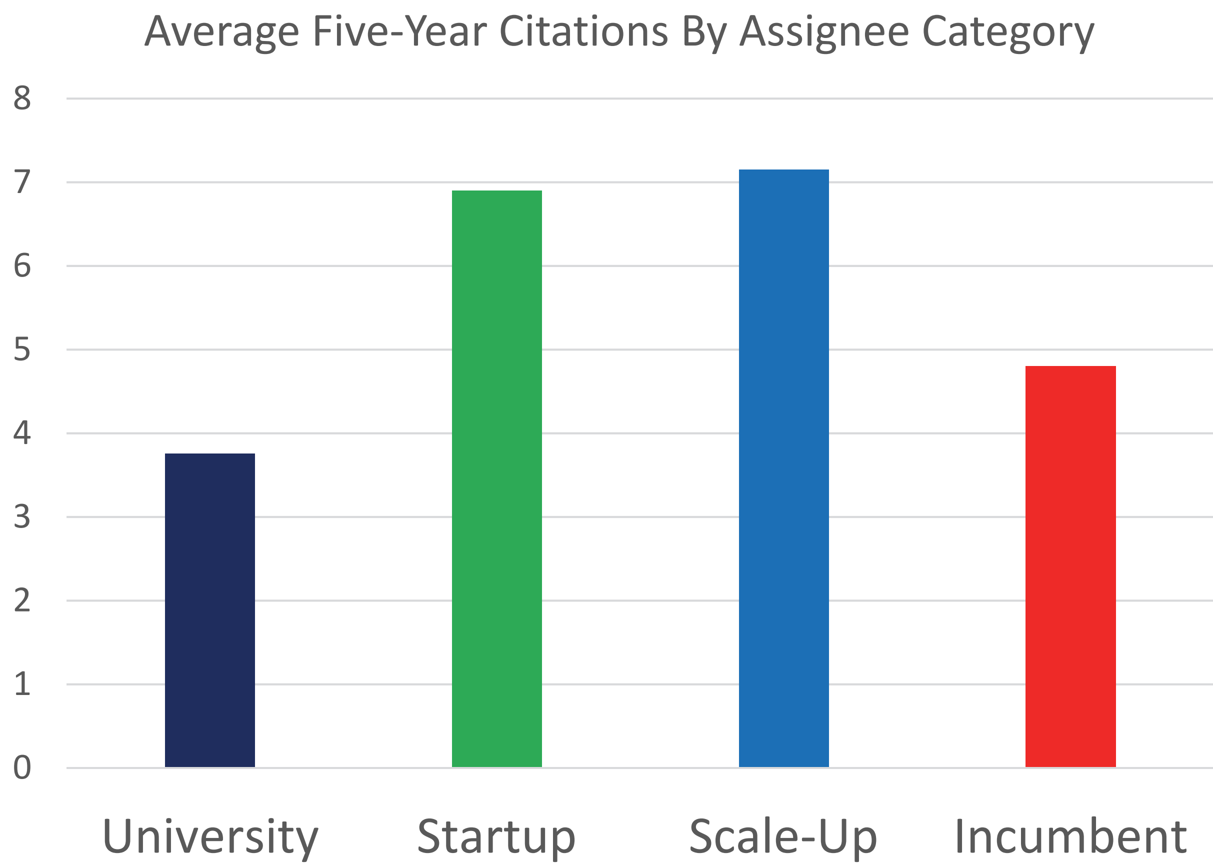

Table A1: Forward Citations

	(1)	(2)	(3)
		DV = Forward Citations	
VARIABLES	0 - 5 Years	6 - 10 Years	11 - 15 Years
Patent Category (Base = Incumbent)			
Startup Patent	1.580***	1.974***	2.444***
	(0.129)	(0.214)	(0.391)
University Patent	1.345***	1.477***	1.578***
	(0.122)	(0.168)	(0.196)
Additional Controls:			
Application and Grant Timing	Y	Y	Y
County Covariates	Y	Y	Y
Grant Year FEs	Y	Y	Y
County FEs	Y	Y	Y
NBER Technology Category FEs	Y	Y	Y
Patent Cohorts:	2000-2005	2000-2005	2000-2005
Observations	27,637	27,637	27,637
Number of Counties	22	22	22
R-squared	0.105	0.128	0.160

Poisson Specifications; Coefficients reported as Incidence Rate Ratios.
Robust standard errors clustered by county in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Table A2: Originality and Generality

	(1)	(2)	(3)
VARIABLES	DV = Originality	DV = 5-Yr Generality	DV = 10-Yr Generality
Patent Category (Base = Incumbent)			
Startup Patent	0.063***	0.057***	0.061***
	(0.016)	(0.011)	(0.013)
University Patent	0.091***	0.078***	0.084***
	(0.015)	(0.011)	(0.016)
Additional Controls:			
Application and Grant Timing	Y	Y	Y
County Covariates	Y	Y	Y
Grant Year FEs	Y	Y	Y
County FEs	Y	Y	Y
NBER Technology Category FEs	Y	Y	Y
Number of Relevant Citations (10 bins)	Y	Y	Y
Patent Cohorts:	2000-2005	2000-2005	2000-2005
Mean of DV	0.425	0.316	0.423
St. Dev. of DV	0.283	0.284	0.277
Observations	26,477	21,949	25,398
Number of Counties	22	22	22
R-squared	0.306	0.392	0.299

OLS Specifications; Robust standard errors clustered by county in parentheses; *** p<0.01, ** p<0.05, * p<0.1