

**Creativity at the Knowledge Frontier:
The Impact of Specialization in Fast- and Slow-paced Domains***

Florenta Teodoridis
USC Marshall School of Business

Michaël Bikard
London Business School

Keyvan Vakili
London Business School

* All authors contributed equally.

Abstract

Using the impact of the Soviet Union's collapse on the performance of theoretical mathematicians as a natural experiment, we attempt to resolve the controversy in prior research on whether specialists or generalists have superior creative performance. While many have highlighted generalists' advantage due to access to a wider set of knowledge components, others have underlined the benefits that specialists can derive from their deep expertise. We argue that this disagreement might be partly driven by the fact that the pace of change in a knowledge domain shapes the relative return from being a specialist or a generalist. We show that generalist scientists performed best when the pace of change was slower and their ability to draw from diverse knowledge domains was an advantage in the field, but specialists gained advantage when the pace of change increased and their deeper expertise allowed them to use new knowledge created at the knowledge frontier. We discuss and test the roles of cognitive mechanisms and of competition for scarce resources. Specifically, we show that specialists became more desirable collaborators when the pace of change was faster, but when the pace of change was slower, generalists were more sought after as collaborators. Overall, our results highlight important trade-offs associated with specialization for creative performance.

Keywords: knowledge specialization and creativity, change and creativity, knowledge frontier, Soviet collapse and productivity of mathematicians

Since Schumpeter (1942), a vast literature has emphasized the crucial role of scientists and inventors as catalysts of economic change. Through their ingenuity, creative workers produce new knowledge and technologies that spur social and economic growth (Romer, 1990), boost or destroy organizational capabilities (Henderson and Clark, 1990), shake old industries (Barnett, 1990), or give birth to new ones (Hargadon and Douglas, 2001). The literature points to two types of creative workers: specialists and generalists. Specialists have experience and deep expertise in a narrowly defined domain of knowledge, while generalists tend to have a large amount of experience but spread across multiple related or unrelated knowledge domains. A generalist with the same amount of experience spread across multiple domains would, by definition, have less expertise in each domain. The distinction is rooted in a strategic trade-off that all creative workers face: either invest their limited time entirely within a specific knowledge domain and become a specialist in that domain or invest it across several domains—achieving a less comprehensive understanding of each—and become a generalist. Past research seems deeply divided on which strategy leads to superior creative performance.

On one hand, a large stream of work has argued that because creativity is about producing novel knowledge recombinations, creative workers should seek access to diverse knowledge bases. Individuals who adopt this approach—generalists—gain access to a wider array of knowledge, technologies, and heuristics that can help them break away from traditional thought patterns (Hargadon and Sutton, 1997; Taylor and Greve, 2006; Fleming, Mingo, and Chen, 2007). Scholars have shown that breakthrough inventions often involve uncommon recombinations of knowledge components from distant domains (Ahuja and Lampert, 2001; Fleming, 2001), and others have found a link between access to atypical knowledge sources and

the creative performance of artists (Cattani and Ferriani, 2008), managers (Burt, 2004), inventors (Reagans and Zuckerman, 2001), and scientists (Schilling and Green, 2011).

On the other hand, several studies have argued for the benefits of specialization because it enables individuals to achieve deeper expertise and a more detailed understanding of the knowledge gaps in their domain of specialty. Specialists also benefit from a more extensive repertoire of domain-specific problem-solving and memory skills (Dane, 2010). Empirical evidence indicates that specialized scientists (Leahey, 2007) and inventors (Conti, Gambardella, and Mariani, 2013) are more productive and successful. Other research has shown that deeper expertise and local recombinations, as opposed to distant and diverse recombinations, are more likely to yield more cognitively novel innovations (Kaplan and Vakili, 2015). Moreover, there is some evidence that individuals lacking specialization might spread themselves too thin and be perceived as less credible (Weisberg, 1998; Leahey, 2007).

We propose that the seeming inconsistency between these two streams of work stems, in part, from efforts to generalize from different settings with different underlying characteristics. The amount of evidence highlighting the superior performance of both specialists and generalists suggests that they have strengths and weaknesses that make them better suited to different circumstances, yet past studies have typically abstracted away from the specificities of their research settings and argued for one view or the other. The respective advantages and disadvantages of generalists and specialists suggest that creative people face a trade-off between strategies that are best suited to different types of circumstances. One of these is the pace at which knowledge components become available in a domain, which can alter the cost-benefit balance of being a generalist or a specialist in different directions. Specifically, we hypothesize that generalists will experience superior creative performance in slower-paced knowledge

domains—access to diverse knowledge enables introducing new knowledge recombinations in these domains—whereas specialists generally benefit from faster-evolving ones because their depth of expertise allows them to exploit new knowledge components faster and more effectively.

Testing this hypothesis empirically is challenging. While the pace of change in a domain may affect the performance of creative workers, their performance also shapes the availability of new components and, with it, the pace of change in that domain (Carnabuci and Bruggeman, 2009). We address this challenge by exploiting a natural experiment—the unexpected acceleration of the pace of change in some fields of theoretical mathematics more than in others after the collapse of the Soviet Union (hereafter the Soviet collapse) in 1989. The event provides a rare opportunity to identify the impact of a change in the pace of knowledge advancement on the relative creative performance of specialist and generalist mathematicians. We investigate how shifts in the pace of change affected the relative creative performance of specialist and generalist mathematicians after the Soviet collapse. In so doing, we aim to reconcile, at least partly, existing debates about the relationship between the strategic choice of becoming a specialist (or generalist) and creative performance.

The Generalist Versus Specialist Trade-off

Creative workers are not born at the frontier of their field. As Amabile (1983: 363) noted, “It is impossible to be creative in nuclear physics unless one knows something (and probably a great deal) about nuclear physics.” While some learning is necessary to contribute to any field, creative workers nevertheless have some discretion about how to allocate their effort. They can choose to invest it narrowly in a specific knowledge domain or seek a broader if shallower

knowledge base. At its core, this choice stems from the underlying tension between an individual's limited time and cognitive abilities and the limitless amount of knowledge that one could potentially acquire (Jones, 2009).

The decision of whether to become a specialist or a generalist is a strategic one.¹ Both are likely to confer distinct strengths and weaknesses, but bounded rationality and limited time mean that one cannot commit to both strategies at the same time. This decision is especially important considering the competitive nature of creative work. Because resources, collaborators, and attention are scarce, small differences in performance can have long-lasting consequences. For example, individuals whose skills are more valued are likely to enjoy easier access to high-quality collaborators, leading to higher future performance. The opposite is also true. Skills that are marginally less valued might prevent the formation of potentially productive collaborations and therefore have long-term negative consequences for creative performance (Merton, 1968; Latour and Woolgar, 1986; Reschke, Azoulay, and Stuart, 2017). Yet prior studies disagree about which creative type is most successful.

Generalists and Creativity

A large stream of literature highlights reasons why generalists are more likely to experience superior levels of creative performance (Nagle and Teodoridis, 2017). Because creative work is a recombination process, individuals who have access to a more varied set of knowledge components can experiment with a larger set of recombinations. Ideas and knowledge from one domain can sometimes be productively applied to others (Hargadon and Sutton, 1997; Burt, 2004; Jeppesen and Lakhani, 2010; Ferguson and Carnabuci, 2017). Moreover, access to diverse knowledge domains is likely to make creative workers more flexible in their problem-solving

¹ Conceptually, generalists and specialists constitute two ends of a continuum. For clarity, our theory considers only these two ideal types.

approaches, allowing them to use a broader repertoire of perspectives and heuristics in their work (Dunbar, 1995).

Becoming a generalist and developing boundary-spanning skills is not costless, however. It is difficult to identify relevant recombinations in general but even more so when searching across knowledge domains. Scientific knowledge tends to be complex, often including tacit components (Polanyi, 1958; Orlikowski, 2002). In any field, the literature is voluminous, and the quality of each contribution is uncertain. Different knowledge communities often use distinct approaches to creating, validating, and describing knowledge (Dougherty, 1992; Knorr-Cetina, 1999). To identify opportunities for recombination across knowledge domains and for implementing them, boundedly rational creative workers must therefore invest considerable time, effort, and resources not only to familiarize themselves with those different approaches but also to build a social network spanning those distinct communities. Those efforts, in addition, often lead to dead ends (Fleming, 2001; Leahey, Beckman, and Stanko, 2017).

Specialists and Creativity

Specialists lack the knowledge breadth of generalists, but their narrow focus allows them to develop a deeper understanding of their domain. Specialists can recall larger amounts of domain-specific knowledge more effectively. For example, Chase and Simon (1973) showed that chess masters can recall the exact position of every piece on a chessboard by observing the board very briefly. In addition, specialists have a more sophisticated appreciation of the different attributes of each component in their knowledge domain as well as of the relationships between those components (Dane, 2010). Chi et al. (1982) showed that expert physicists tend to categorize domain-specific physics problems according to physics principles, whereas non-experts categorize according to features noted in the problem statement. In turn, categorization based on

principles helps experts activate knowledge structures related to each principle, helping problem-solving activities.

These benefits notwithstanding, specialization also has important downsides. It might lead to situations in which the same set of familiar components is constantly used and reused, leading to decreasing returns to creative work (Fleming, 2001). Moreover, psychologists found that knowledge specialization leads to the development and reinforcement of thought processes that become taken for granted, a phenomenon known as the “Einstellung” effect (or “problem-solving fixation”) (Luchins, 1942; Frensch and Sternberg, 1989; Bilalić, McLeod, and Gobet, 2008: 653). The effect was famously documented in Abraham Luchins’ (1942) water-jug experiment. In this short study, participants faced a set of five problems—“Einstellung Problems”—all of which could be solved in the same manner. Following this, they were given another set of problems that could be solved laboriously with this method but for which a much simpler method also existed. Of the participants exposed to the full set of Einstellung Problems, none found the simpler solution. In contrast, over 60 percent of participants in the control group who were warned to be careful after their exposition to the Einstellung problems identified the simpler solution. In other words, prior experience led to routinized problem solving which in turn blinded participants to the existence of a better solution. This striking demonstration of the negative impact of expertise on creativity has since been replicated and extended in a large number of studies (see Bilalić, McLeod, and Gobet, 2008; Dane, 2010).

Specialization in a domain can therefore lead to the formation of habitual behaviors rooted in one’s knowledge structure (Aarts, Verplanken, and van Knippenberg, 1998; Aarts and Dijksterhuis, 2000; Audia and Goncalo, 2007; Murray and Häubl, 2007; Chai, 2017). While

habits can increase efficiency in dealing with the routine tasks required in a domain, they can nonetheless slow the pace of adaptation to new tasks and heuristics.

The Pace of Change as a Moderator of Creative Performance

Clearly, the decision to become a specialist or a generalist implies different advantages and drawbacks. Thus some studies have attempted to reconcile the seemingly inconsistent findings in the literature on the superiority of either strategy by highlighting differences in the quantity and quality of creative output. The general idea is that the generalists' access to more diverse knowledge components enables them to produce more novel recombinations, but while some of these novel recombinations may have great impact, many others will fail. In contrast, specialists face fewer barriers to recombination, but they lack access to the same variety of components. One might therefore expect them to produce more incremental recombinations (Carnabuci and Bruggeman, 2009; Leahey, Beckman, and Stanko, 2017).

We take a different approach to reconcile the inconsistency in prior research, which for the most part has neglected the underlying characteristics of the context within which creative workers operate. Previous studies have either explored the average performance of a certain creative strategy in an aggregated sample of creative workers across many different domains (e.g., Fleming, Mingo, and Chen, 2007) or investigated the performance of creative workers in a single domain of knowledge or technology (e.g., Audia and Goncalo, 2007). Despite the usual cautionary notes about the generalizability of findings, both categories of research have generally overlooked how the variance in characteristics of knowledge domains moderates the performance of different creative strategies. Our aim here is to take a step toward addressing this gap by focusing on a specific characteristic of a knowledge domain, namely the pace at which

new knowledge components become available in the domain, and by showing how it differentially affects the creative performance of specialists and generalists.

We build our argument on three premises. First, innovation is most often a process of knowledge recombination (Fleming, 2001). New ideas are essentially combinations of previously unconnected knowledge components. Creativity is cumulative, and each new recombination is a new component that can be used for future discoveries or inventions (Weitzman, 1998).

Second, the set of knowledge components in a domain is not fixed. As scientific and technological advancements occur in a domain, new knowledge components become available. The emergence of new knowledge components thus affects the set of opportunities available to creative workers for knowledge recombination. The more knowledge components emerge in a domain, the more opportunities emerge for recombinations between the new knowledge components themselves, as well as between the new knowledge components and the previously established ones.

Third, the pace of change varies substantially across knowledge domains and over time. Periods of intense change often alternate with more stable periods (Kuhn, 1970; Dosi, 1982) in a process analogous to the punctuated equilibrium framework developed by evolutionary biologists (Gould and Eldredge, 1977). Numerous episodes of sudden acceleration or deceleration have been described. Tushman and Anderson (1986) showed that the cement, airline, and minicomputer industries experienced periods of rapid technological change followed by years of relatively slow improvements. Lim (2009) documented how IBM's breakthrough development of copper interconnects to replace aluminum ones in 1999 moved the knowledge frontier in the semiconductor industry considerably and paved the way for the production of smaller chips with superior conductivity. The biotechnology industry in its early stages was

reportedly shaped and shaken by various scientific discoveries in genetics (Russo, 2003). Levinthal (1998) described how the broadcasting industry emerged almost overnight after the abrupt realization that there was a demand for this type of technology. In science, Boring (1955) noted that the invention of the telescope in 1608 enabled a flurry of astronomical discoveries. Kuhn (1970) described how the realization by Joseph Black in 1756 that air was not the only gas—and his identification of fixed air (CO₂)—opened the door to the rapid discovery of numerous other gases by Cavendish, Priestly, and Scheele. More recently, the Human Genome Project is credited with the birth of the new field of genomics and a nearly exponential increase in the pace of disease gene discovery.²

These three premises suggest that creative recombinations do not take place only “horizontally”—within or across domains—but also “vertically”—using components that are more or less distant from the frontier. Moreover, the pace of change in a domain is likely to shape the returns on attempting these two types of recombination. A faster pace of change facilitates vertical recombination because it ensures that many new components are available at the frontier. In contrast, a slower pace of change makes vertical recombination more difficult, therefore presumably increasing the relative value of horizontal recombination. While the pace at which new knowledge components become available in a domain influences all knowledge workers in that domain, we expect the pace to have differential effects on the performance of specialists and generalists. In the case of a slow-paced knowledge domain, where the set of knowledge components available for recombination is relatively stable, specialists, who largely rely on the knowledge available within the domain, gradually exhaust any novel and impactful recombinations (Fleming, 2001; Schilling, 2005), a situation that Fleming (2002) called

² <https://www.genome.gov/10001772/all-about-the--human-genome-project-hgp/>.

“combinatoric exhaustion.” Specialists see the benefits of specialization erode while they struggle to adopt fresh perspectives by borrowing ideas from other domains. Slower-paced environments should therefore benefit generalists. Their access to several knowledge domains opens the door to a greater variety of potential recombinations, and the relatively slow emergence of new knowledge components within the field gives them enough time to avoid falling behind their specialist colleagues.

The situation changes when the pace of change in a field accelerates. Specialists are generally better equipped than generalists to take advantage of a faster-evolving knowledge frontier. Their narrow focus means that they can invest their time effectively to absorb newly emerging knowledge components. They can take advantage of their deeper understanding of the field not only to identify knowledge gaps but also to gauge recombination opportunities between new components and those previously available. The drawbacks of specialization also become less important because the rapid emergence of new components within the field lowers the value of borrowing knowledge components from other areas. Given the closer similarity and relevance of newly emerged knowledge components to extant knowledge components in a domain, the recombinations of the two could be less prone to failure than recombinations of components across different knowledge domains. In contrast, generalists in faster-paced environments are likely to struggle to maintain their ties to various fields while staying current with the advancements of the faster-paced domain. Their reliance on riskier boundary-spanning recombinations becomes a liability. With an increase in the pace of change in a domain, specialists’ disadvantage relative to generalists should gradually decline and their relative advantage gradually increase such that the creative performance of specialists may surpass that of generalists in sufficiently fast-paced domains. Therefore we expect:

Hypothesis 1 (H1): As the pace of change in a domain accelerates, the creative performance of specialists relative to generalists increases.

To this point, we have focused on the distinctive capabilities of specialists and generalists. Creative performance, however, is not determined by creative capabilities alone. The ease in producing creative recombinations varies over time (Kuhn, 1970; Dosi, 1982). An acceleration of the pace of change might offer an abundance of new knowledge components that creative workers can productively recombine. It might also attract additional attention and resources for the fast-changing field. As a result, an increase in the pace of change can potentially provide more recombination opportunities for all creative workers in a domain. Hence, even though it advantages specialists more than generalists, one might expect that generalists would benefit, too, but it is more likely that generalists' creative performance will decline when the pace of change increases in a domain, for two reasons.

First, generalists' skills are not well adapted to take advantage of recombination opportunities in a faster-paced domain. They might struggle to identify these opportunities because of their more superficial understanding of the domain and their divided efforts to keep up to date across multiple domains. Second, even if generalists identify some of the emerging opportunities, they are at a disadvantage in attracting the complementary resources needed to act on those opportunities relative to specialists. Creative workers frequently stumble upon similar ideas, but only those that win the priority race can reap the reward for their investment (Merton, 1957). In science, for example, a researcher's success in attracting citations and credit often means that their peers and competitors will receive fewer citations and overall less credit (Reschke, Azoulay, and Stuart, 2017). Complementary resources such as funding, equipment, and collaborators are scarce and will naturally flow to those who are expected to be more

successful (Merton, 1968; Latour and Woolgar, 1986). Even if generalists are able to attract the complementary resources they need, they might still struggle. Creative insights are only valuable to the extent that others learn about them, but attention is scarce. At one extreme, if the work of specialists attracts all the limelight, the work of generalists will remain unknown. We therefore hypothesize:

Hypothesis 2 (H2): As the pace of change in a domain accelerates, generalists are likely to experience a decline in their creative performance; the reverse will occur for specialists.

Methods

Empirical Setting

To test these predictions, we focused on the field of theoretical mathematics and the publication output of mathematics scientists. We follow a growing literature using scientific publications to measure scientists' creative output (e.g., Jones and Weinberg, 2011; Uzzi et al., 2013; Leahey, Beckman, and Stanko, 2017). Moreover, we exploited a natural experiment—the Soviet collapse, in 1989—to address the endogeneity issues involved with testing our predictions. For several reasons, this event provides a unique opportunity to examine the relative performance of specialists versus generalists in knowledge domains with varying paces of change.

The unexpected, exogenous release of new knowledge in certain areas of theoretical mathematics due to the Soviet collapse enabled us to control for the endogenous link between the activity of creative workers and the pace of change of knowledge domains. Our empirical strategy relies on the assertion that the Soviet collapse caused a sudden and unexpected increase in the pace of change in theoretical mathematics and that it did so more for some subfields of mathematics than for others (Agrawal, Goldfarb, and Teodoridis, 2016). We based this claim on three main observations. First, the Soviet Union was, and Russia continues to be, a world-

renowned center of scientific research, with mathematics holding a prominent position. Scholarly research in theoretical mathematics attracted great minds, as it was uniquely detached from politics, conferred status and prestige, and offered financial rewards superior to those of many other occupations. Second, although Soviet mathematics was strong across the entire spectrum of mathematics, Soviet mathematicians made greater advancements in some subfields than in others (Graham, 1993). These differences reflect historical path dependency. Some subfields of theoretical mathematics built on strong mentorship from the early 1900s and continued to attract bright minds thereafter (Borjas and Doran, 2012). For example, the success of Moscow mathematics can be traced back to Ergorov and his student N. N. Luzin (Tikhomirov, 2007), whose famous work focused mainly on the theory of functions. Finally, Soviet knowledge in theoretical mathematics was kept secret from the outside world because of the Communist government's rules and regulations. The Soviet government strictly controlled international travel. Academics seeking to attend foreign conferences had to undergo a stringent and lengthy approval process, and many researchers were blacklisted because of their "tainted" backgrounds. The few approvals granted were typically for travel in Eastern Europe (Ganguli, 2014). Additionally, Soviet researchers were prevented from publishing their findings outside the Soviet Union, from communicating or collaborating with non-Soviets, and even from accessing non-Soviet references. Thus Soviet advancements in mathematics remained relatively unknown to the outside world until the Soviet collapse (Graham and Dezhina, 2008), when they were suddenly made available.³

³ The following quote, from an article published on May 8, 1990, in the *New York Times*, indicates the sudden outward shift of the knowledge frontier: "Persi Diaconis, a mathematician at Harvard, said: 'It's been fantastic. You just have a totally fresh set of insights and results.' Dr. Diaconis said he recently asked Dr. Reshetikhin for help with a problem that had stumped him for 20 years. 'I had asked everyone in America who had any chance of knowing' how to solve a problem of determining how organized sets become disorganized, Dr. Diaconis said. No one could help. But Dr. Reshetikhin told Dr. Diaconis that Soviet scientists had done a lot of work on such problems. 'It was a whole new world I had access to,' Dr. Diaconis said."

Using an extensive dataset of publication and citation data in the field of mathematics, we carefully tracked the creative output and performance of mathematicians over a long period (1980–2000). The data come from the Mathematical Reviews (MR) division of the American Mathematical Society (AMS). The MR Database includes all academic publications in mathematics worldwide.

We observed the specialization levels of mathematicians in our sample based on a manual, detailed categorization of research output provided by the MR Database, which classifies each paper in mathematics using Mathematics Subject Classification (MSC) codes. The MSC schema are internationally recognized and facilitate targeted searches on research subjects across all subfields of mathematics. The MR team assigns one primary MSC code to each academic publication uploaded to the MR Database. There are 33 codes covering theoretical mathematics, as described below. Using the MSC codes assigned to each paper, we can measure the degree of specialization of each individual mathematician at a given time.

The field of theoretical mathematics plays a fundamental role in knowledge and technological progress across a wide range of domains. Wavelet and Fourier transforms are widely used in electronics, computer graphics, and medical equipment such as MRI machines. Algebraic topology is used extensively in data mining and processing. Number theory, particularly the theory of prime numbers, has immensely influenced computer and network security algorithms. Turing's theories of computability provided the foundation for the field of computing. Many advancements in space technology and exploration would have been impossible without foundational geometry theories. Theoretical math has substantially influenced many areas in the social sciences such as linguistics, economics, and political science. Put simply, theoretical mathematics provides the abstract foundation and structure for

formulating and understanding our physical world. Corporations such as Microsoft, Google, and IBM employ theoretical mathematicians in various areas of security and computing. Hence the field of theoretical mathematics provides valuable insights into one of the fundamental engines of economic, technological, and social progress.

Data

The MR Database covers the three main branches of mathematics: mathematical foundations (including history and biography), pure or theoretical mathematics, and applied mathematics. Our focus is on theoretical mathematics, which includes analysis, algebra, and geometry. Our sample tracks academic publications of mathematicians over a 21-year period, 1980 to 2000 inclusive.

To construct our sample, we first collected data on every academic publication in theoretical mathematics published between 1980 and 2000, 10 years before and after the collapse of the Soviet Union in 1989. The data on publications include year of publication, MSC classification code, full set of authors per academic publication, and number of academic citations received from subsequent publications. Next, we rearranged the data at the author-year level and counted the number of academic publications and citations per author, per year. We excluded all Soviet authors, who were already at the frontier of knowledge, and focused on all other mathematicians, who experienced the frontier advancement. We also excluded all publications with at least one Soviet author to ensure that our results are not driven by preferential direct access to Soviet knowledge. We further restricted our sample to authors with at least four publications before the Soviet collapse, namely between 1980 and 1989.⁴ The choice of a minimum of four publications helped us carefully separate specialists from generalists in our

⁴ The results are robust to choosing cut-off minimums of three, five, and six publications. At cut-offs smaller than three publications, we cannot properly distinguish between specialist and novice mathematicians.

sample and ensure that our results are not driven by unproductive individuals classified as specialists because of their low number of publications. For example, individuals with one publication would otherwise be automatically classified as specialists, but their lack of diversification would be mechanically driven by their low productivity. We provide details on our measure of diversification in the next sections. Finally, using the diversification measure described below, we identified all individuals who could be cleanly categorized as either a specialist or a generalist and dropped the rest from the sample. In our robustness checks, we provide sensitivity analyses on our categorization of specialists and generalists. The final core dataset contains data on 6,358 mathematicians and their full record of publications between 1980 and 2000.

Last, we matched specialists and generalists on their productivity in the period before the Soviet collapse. As we discuss below, there are some significant differences in productivity between specialists and generalists in the years before the collapse. This is not surprising, because our measure of diversification relies on breadth of publications across mathematics subfields. In other words, the higher the productivity, the higher the probability of diversification. Thus to ensure that our results are not biased because of systematic differences in quality between specialists and generalists driven by our sample-selection method, we further constructed a matched sample based on individuals' observables before the Soviet collapse. To construct the matched sample, we used a one-to-one coarsened exact matching (CEM) method (Blackwell et al., 2009; Iacus et al., 2011) based on mathematicians' publication records in the pre-collapse period.⁵ The matched sample contains data on 4,076 mathematicians, of whom

⁵ To perform the one-to-one matching, we used the total citation-weighted number of publications, total number of publications, the first year of publication, and the publication trend during the 10 years prior to the Soviet collapse.

2,038 are specialists and 2,038 are generalists. We report our estimations for both the full and the matched samples.

Dependent Variables

To compare the creative output of specialists and generalists, we used three variables to capture both the quantity and the quality of their output. First, we used the count of publications per year to measure the quantity of their creative output. The issue with using the simple count of publications is that an increase in the number of publications may come at the expense of a decrease in their quality. To address this issue, we also used the quality-adjusted count of publications per year. Following previous studies (e.g., Furman and Stern, 2011; Azoulay, Stuart, and Wang, 2013; Vakili and McGahan, 2016), we used the number of citations each publication received in subsequent publications by 2014 to construct a citation-weighted count of publications per mathematician per year.⁶ Each publication is counted as 1 plus the number of future citations it received. For example, if a mathematician had two publications in 1985, one with 10 future citations and the other with 20 future citations, their quality-adjusted research output for 1985 is 32. We also measured the number of breakthrough publications per year for each scientist. The quality of creative output is highly skewed. Past research distinguished between processes that increase the mean distribution of creative output and those that increase the variance. While the former can raise the average quality of creative output, the latter can lead to an increase in the number of highly impactful publications—that is, breakthroughs. Following past research (Ahuja and Lampert, 2001; Phene, Fladmoe-Lindquist, and Marsh, 2006; Bikard, Murray, and Gans, 2015; Kaplan and Vakili, 2015), we first coded the publications belonging to the top 5 percent of highly cited publications in any given year as breakthroughs. Next, we

⁶ While older publications have more time to collect citations, the inclusion of year fixed effects in our regressions ensures that each publication is only compared with other publications in the same year.

counted the number of breakthroughs for each individual mathematician in any given year to construct an individual measure of breakthrough output per year. As a robustness check, we also constructed a separate measure of breakthrough output per year based on publications in the top 10 percent of highly cited publications.

Independent Variables

We used three indicators (and their interactions) as main independent variables in all estimations. The first variable, $Specialized_{i,t}$, captures whether a mathematician in our sample is a specialist or a generalist at the time of the Soviet collapse. To construct this variable, we first built an index of diversification at the individual level capturing the heterogeneity in breadth of knowledge based on each mathematician's publication portfolio during the period before the collapse (1980–1989). The index is calculated as 1 minus the Euclidian distance in the multidimensional space of 33 subfields (or MSC codes) of theoretical mathematics and is based on shares of publications in each of the 33 subfields, per mathematician. The Euclidian distance is equal to the square root of the Herfindahl index and hence is a more conservative measure of diversification. Formally, we calculated:

$$DiversificationIndex_i = 1 - \sqrt{\sum_{s=1}^{33} \left(\frac{PubCount_{s,i}}{PubCount_i}\right)^2}$$

By construction, the higher the value of $DiversificationIndex_i$, the greater the breadth of areas in which mathematician i published before the Soviet collapse. The diversification measure is greater than or equal to 0 and never reaches 1. The highest possible value of the diversification index is .83 and characterizes researchers who published an equal number of publications in all 33 subfields of theoretical mathematics. The lowest diversification index is 0 and characterizes mathematicians who published in one subfield of theoretical mathematics exclusively. For example, a mathematician who published a total of 10 papers, half in one subfield of theoretical mathematics and half in another, would have a diversification index of .29, and an equally productive colleague who published all their papers in one subfield of theoretical mathematics would have a diversification index of 0. In our sample, the highest diversification index is .531 and the lowest is 0. In our main specification, we define generalists as mathematicians having a diversification index in the top 10 percent of the distribution (above .290) and specialists as those

having a diversification index of 0 (those who published in only one subfield). Our results remain robust to using a continuous measure of diversification. See tables A1 to A3 in the Online Appendix for robustness checks (<http://journals.sagepub.com/doi/suppl/10.1177/0001839218793384>).

[Insert table 1 about here]

The second variable, *SovietImpact_i*, captures the degree to which each mathematician in our sample was affected by the Soviet collapse or, in other words, by the pace of advancement of the knowledge area. The variable separates mathematicians who experienced a substantial movement of the knowledge frontier in their areas, that is, those operating in a fast-paced knowledge domain, from mathematicians who experienced less of a movement, that is, those operating in a slow-paced domain. We followed the ranking in Agrawal, Goldfarb, and Teodoridis (2016) of the 33 primary MSC codes of theoretical mathematics indicating the degree to which Soviets contributed to each subfield before the Soviet collapse. Table 1 lists the 33 subfields and their ranks. Based on these rankings, we constructed an index of Soviet exposure for each scientist in our dataset who published between 1980 and 1989. The index is calculated as the sum of shares of publications in each of the 33 subfields of theoretical mathematics, weighted by the ranking of the 33 subfields, per individual, for the entire period before the Soviet collapse. The higher the percentage of one’s publications in subfields where Soviets made greater contributions, the higher the Soviet impact index. Formally, we calculated:

$$SovietImpactIndex_i = \sum_{s=1}^{33} \frac{PubCount_{si}}{PubCount_i * SubfieldRankOrder_s}$$

where *PubCount_{si}* is the total count of publications of scientist *i* in subfield *s*, *PubCount_i* is the total count of publications of scientist *i*, and *SubfieldRankOrder_s* is the rank order of the corresponding subfield *s* in theoretical mathematics. The calculation considers the full publication portfolio during the period before the collapse (1980–1989). For example, a

mathematician who published all their papers in “Integral Equations,” the most affected subfield of theoretical mathematics, would have a Soviet impact index of 1. If that person were to publish all their papers in “Fourier Analysis,” the second most affected subfield of theoretical mathematics, their Soviet impact index would be .5. And if that person were to publish half their work in “Integral Equations” and half in “Fourier Analysis,” their Soviet index impact would be .75. In our sample, the minimum value of the Soviet impact index is .030, the maximum value is 1, the mean is .108, and the standard deviation is .112. We defined mathematicians most affected by the Soviet shock ($SovietImpact_i = 1$) as those having a Soviet impact index in the top 10 percent of the range. The indicator is equal to 0 for others. Our results remain robust to considering a continuous measure of Soviet impact. See tables A4 to A6 in the Online Appendix for robustness checks.⁷

The third variable, $AfterSovietCollapse_i$, is an indicator equal to 1 for years after the collapse of the Soviet Union (1990 and after) and 0 otherwise.

Control Variables

In all estimations, we included individual and year fixed effects. Individual fixed effects controlled for all time-invariant, idiosyncratic characteristics of each mathematician, such as first year of publication, innate quality, gender, race, and year of graduation. The year fixed effects controlled for all macro time trends that could influence mathematicians in the sample.

We also controlled for the past productivity of mathematicians using the cumulative number of publications (since 1980). The variable is logged to account for its skewed distribution. Last, we controlled for the nonlinear effect of age by including an age-squared term

⁷ We use a 0/1 indicator instead of a continuous variable for ease of exposition. Our estimations rely on a triple interaction between our independent variables; hence using 0/1 indicators facilitates interpretation of the magnitude of the estimation results.

in all estimations.⁸ Because we could not observe the actual age of individuals, we used the number of years since their first publication in our sample.

Estimation Strategy

We used a difference-in-difference-in-differences (DDD) estimation method to compare the research output of specialists and generalists affected by the forward movement of the knowledge frontier in theoretical mathematics due to the Soviet collapse. The DDD estimation strategy is meant to address the endogeneity of output behavior and forward movement of the frontier by controlling for the underlying difference in the performance of specialists and generalists in relation to the forward movement of the frontier. Formally, we estimated:

$$DV_{i,t} = f(\beta_1 \cdot Specialist_i \cdot SovietImpact_i \cdot AfterSovietCollapse_t + \beta_2 \cdot Specialist_i \cdot AfterSovietCollapse_t + \beta_3 \cdot SovietImpact_i \cdot AfterSovietCollapse_t + C_{i,t} + I_i + \gamma_t + \varepsilon_{i,t})$$

where $DV_{i,t}$ represents mathematician i 's output of interest (citation-weighted publication count, breakthrough count, and collaboration rate) in year t . $Specialist_i$, $SovietImpact_i$, and $AfterSovietCollapse_t$ are the three main independent variables. $C_{i,t}$ represents the set of control variables (cumulative number of publications and age squared). I_i and γ_t indicate individual and year fixed effects, respectively. Note that $Specialist_i$ and $SovietImpact_i$ are not included independently because they are absorbed by individual fixed effects, as their values are fixed at the individual level. Similarly, $AfterSovietCollapse_t$ is not included independently because its effect is absorbed by the year fixed effects.

β_2 captures the difference between post-Soviet outcome trends of specialists and generalists whose research was primarily in slow-paced areas—areas less affected by the Soviet

⁸ The inclusion of both individual fixed effects and year fixed effects automatically controls for the linear effect of individual age.

collapse. β_1 is the main coefficient of interest for testing H1 and captures the differential performance of specialists relative to generalists in fast-paced areas of theoretical mathematics—areas most affected by the Soviet collapse—using the difference in areas less affected as the baseline. β_3 captures the change in the outcome trend of generalists who were active in fast-paced areas compared with generalists whose research was predominantly in slow-paced areas, and $\beta_1 + \beta_3$ captures the equivalent change between specialists in fast- and slow-paced areas. Together, β_3 and $\beta_1 + \beta_3$ are the coefficients of interest for testing H2.

Because all three dependent variables are count variables, we used a conditional fixed-effects panel Poisson model with robust standard errors clustered at the individual level and calculated using the Huber–White method in all estimations. The estimator is consistent in the presence of heteroskedasticity and overdispersion of the dependent variable (Silva and Tenreyro 2006).

Results

[Insert table 2 about here]

Descriptive statistics and correlations for the full sample and the matched sample are shown in table 2. In the full sample, a typical mathematician in our sample has produced approximately .7 papers per year (or about two papers every three years) and has a citation-weighted publication count of approximately 5.7. She has also produced, on average, one publication in the top 5 percent and two publications in the top 10 percent during the whole sample period (1980–2000). Note that the figures are skewed. Hence while many mathematicians in our sample have not produced any breakthroughs, others have produced multiple breakthroughs during the sample period. Furthermore, a typical mathematician has collaborated with at least one person every

other year, and most of her collaborations are unique. The means for the matched sample are slightly smaller than those for the full sample due to a lack of proper matches for individuals with extremely high levels of productivity. Nevertheless, overall there is substantial overlap between the full sample and the matched sample.

[Insert table 3 about here]

Table 3 details the differences between specialists and generalists on the key dimensions of interest for the period before the Soviet collapse. Panel A shows the differences in the full sample, and panel B reports them in the matched sample. In the full sample (panel A) there are almost twice as many specialists as generalists. This is not surprising given their graduate training and the importance of establishing a domain of specialty for career advancement in academia (e.g., Franzoni, Scellato, and Stephan, 2011; Stephan, 2012). The generalists in our sample produce on average approximately one more publication and 17 more citation-weighted publications in the period before the collapse. Furthermore, generalists generate approximately .2 more publications in the top 5 percent cited list, relative to specialists, in that period. In other words, a typical generalist is 1.6 times more likely to produce highly cited publications. The difference is similar when focusing on the number of publications in the top 10 percent cited. Interestingly, specialists seem to collaborate more frequently on their papers but have fewer unique collaborators. This is consistent with the idea that generalists are more likely to work with a more diverse set of individuals across a wider range of domains. Panel B presents the comparative descriptive statistics for the matched sample. The main takeaway is that the differences between specialists and generalists in the full sample disappear once we restrict our sample to the CEM one. Last, due to our strict one-to-one matching, the numbers of specialists

and generalists are the same in the matched sample. The number of generalists in the matched sample is not considerably different from the number of generalists in the full sample.

[Figure 1 about here]

Figure 1 shows the difference between the citation-weighted number of publications produced by specialist and generalist mathematicians after the collapse of the Soviet Union across the faster-paced areas (i.e., most affected by the Soviet collapse) and slower-paced areas (i.e., least affected by the Soviet collapse). Figure 1a is based on the full sample, and figure 1b is based on the matched sample. In both graphs, the post-collapse values are adjusted based on the pre-collapse values. Specifically, for each group of mathematicians (specialists in faster-paced areas, specialists in slower-paced areas, generalists in faster-paced areas, and generalists in slower-paced areas), we display their weighted post-collapse output. The weighting is based on the ratio of their pre-collapse output to that of specialists in the slower-paced areas.⁹ This approach allows us to offer a concise visual representation of differences in creative output of all four groups of mathematicians post collapse.

The graphs suggest that specialist mathematicians in faster-paced areas produced more citation-weighted publications than generalist mathematicians in those areas. In contrast, in the slower-paced areas, specialist mathematicians produced significantly fewer citation-weighted publications than generalist mathematicians. The observed differences are consistent with hypothesis 1. Moreover, generalist mathematicians in faster-paced areas produced fewer citation-weighted publications than generalists in slower-paced areas, consistent with hypothesis 2.¹⁰

⁹ The approach is neutral to the choice of baseline group.

¹⁰The graphs also suggest that generalists in faster-paced domains do worse than specialists in slower-paced domains. This seems to indicate that the downsides of being a generalist are worse than those of being a specialist, but this result should be interpreted carefully. Our theory suggests that as the pace of change increases, specialists thrive and generalists suffer. Hence it is possible that the difference between specialists in slower and generalists in faster domains is driven by the fact that the “slower” domains we observe are in fact not very slow, whereas the “faster” domains are indeed very fast. In other words, the observation comparing generalists and specialists across

[Table 4 about here]

Next we tested these predictions using regression analysis. Table 4 shows estimation results for the change in the creative output of specialists and generalists in slower- and faster-paced domains. Model 1 of table 4 shows estimation results using count of publications as the dependent variable. The estimated β_2 suggests a statistically significant 8-percent relative decline in the number of publications by specialists in slow-paced areas of theoretical mathematics compared with generalists in those areas.¹¹ The decline is equivalent to approximately one fewer publication by specialists compared with generalists after the collapse. In contrast, there is a relative increase of approximately 37 percent in the number of publications by specialists over generalists in faster-paced areas of mathematics, using the change in the differential performance of specialists relative to generalists in less affected areas as the baseline. The 37-percent increase is equivalent to approximately three extra publications after the collapse. The results are consistent with H1, suggesting that specialists have higher creative performance than generalists in faster-paced areas. The negative and significant β_3 suggests a 23-percent decrease in creative output of generalists in faster-paced knowledge domains compared with generalists in slower-paced knowledge domains, consistent with H2. Also, although not statistically significant ($p = .35$), compared with specialists in the less affected areas of mathematics, specialists in the most affected areas increased their publication count by approximately 5 percent after the Soviet collapse (based on the sum of β_1 and β_3).

slower- and faster-paced domains is arguably context-specific and depends on the level of increase in the pace of change in a domain and the relative length of the period during which the domain experiences faster or slower changes.

¹¹ To calculate percentage change in output trends, we compute the incidence rate ratio from the estimated coefficients.

Model 2 shows results for the citation-weighted number of publications. The interpretation of results is similar to those reported for model 1. The coefficients are larger, however, which suggests that the relative increase in creative performance of specialists in affected areas after the Soviet collapse is driven partly by an increase in the quantity of their creative output and partly by an increase in the average quality of their creative output (measured as the number of citations to their publications). The estimated β_2 suggests that, in slower-paced areas of mathematics, specialists produced approximately 22 percent fewer citation-weighted publications per year than generalists did in years after 1989. The decline is equivalent to producing approximately three fewer citation-weighted publications per year after the collapse. In contrast, when we use the change in the differential performance of specialists versus generalists in less affected areas as the baseline, the estimated β_1 suggests that in faster-paced areas of mathematics, specialists increased their citation-weighted publication output relative to generalists by approximately 83 percent in years after 1989. This is equivalent to producing approximately four more citation-weighted publications per year during the post-Soviet period. The negative and significant β_3 suggests a 37-percent decrease in creative output of generalists in faster-paced knowledge domains compared with generalists in slower-paced knowledge domains. The results also indicate that, compared with specialists in less affected areas, specialists in the most affected areas increased their performance by a statistically significant margin of approximately 16 percent after the Soviet collapse.

Models 3 and 4 report the analog estimation results for the matched sample. The estimated coefficients of β_1 and β_3 are slightly larger. They are in line with those for the full sample and depict trends aligned with those described above. Overall, the results in table 4 provide strong support for the effects hypothesized in H1 and H2.

One potential concern with these interpretations is that the observed change in the differential performance of specialists and generalists in the faster- and slower-paced areas of mathematics might have begun before the Soviet collapse and that our estimations are driven by these pre-trends. To address this concern, we checked the timing of changes in specialists' performance by estimating their differential performance relative to generalists' in years before and after the Soviet collapse in 1989. We used the 1987 to 1989 performance difference of specialists and generalists in the non-affected areas as the baseline (i.e., the time right before the Soviet collapse) and examined the change in citation-weighted output of specialists and generalists over six periods: 1981–1983, 1984–1986, 1990–1992, 1993–1995, 1996–1998, and 1999–2000. The use of citation-weighted output helps us capture the changes in both quantity and quality over time. The estimations are based on the same DDD estimator as before, where we replace the *AfterSovietCollapse_t* dummy with dummies for each of the three-year periods described. We used groups of three years because many mathematicians in our sample publish once every few years. In testing for pre-trends, we used the unmatched sample (i.e., full sample) to ensure that the pre-trends are not masked by our matching procedure. In the Online Appendix, we show that the graphs are similar if we use the matched sample (figure A1). If changes in the differential performance of specialists and generalists indeed predate the Soviet collapse, we should observe these trends in the 1981–1983, 1984–1986, and 1996–1998 periods. In figure 2a we plot the estimated yearly β_2 coefficients, which represent the difference in citation-weighted outputs of specialists relative to generalists in slower-paced areas. We observe a decrease in the estimated difference in output between specialists and generalists in slower-paced areas after the collapse of the Soviet Union.

[Insert figure 2 about here]

In figure 2b we plot the estimated yearly β_1 , which represents the relative difference in citation-weighted outputs of specialists relative to generalists in faster-paced areas. The estimates suggest that the change in the differential performance of specialists and generalists in these faster-paced areas increases after the fall of the Soviet Union, in line with H1. In figure 2c we plot the estimated yearly β_3 , which represents the difference in citation-weighted outputs between generalists in faster-paced areas and generalists in slow-paced areas. Figure 2d further shows the difference in citation-weighted output between specialists in faster-paced areas and specialists in slower-paced areas ($\beta_1 + \beta_3$). In line with H2, we observe a decrease in the performance of generalists in faster- versus slower-paced areas of theoretical mathematics after the Soviet collapse. In contrast, we observe an increase in the performance of specialists in faster-paced areas after the collapse. In all figures there are no indications of pre-Soviet collapse trends, confirming that the estimated changes in table 1 are attributable to years following the fall of the Soviet Union.

In table A7 in the Online Appendix we examine the differential propensity of generalists and specialists to produce breakthroughs in faster- and slower-paced knowledge domains. The estimates are in line with those reported in table 1 for overall creative output and are consistent with H1 and H2. Theoretically, the findings imply that specialists in faster-paced areas not only absorb the newly emerged knowledge components faster than generalists but also use the new knowledge more quickly to address the more fundamental gaps in their domain of specialty. Their faster and more effective absorption and use of new knowledge can potentially crowd out generalists' efforts and push generalists to tackle less impactful opportunities. The opposite argument holds for generalists in slower-paced domains. The results suggest that the usual trade-

off between quantity and quality does not necessarily hold in all contexts. In certain conditions, one creative type may show superior performance on both dimensions of quantity and quality.

We also tested the role of competition for complementary resources as a force driving the differential performance of specialists and generalists after the Soviet collapse. We examined this mechanism by investigating the change in the collaboration patterns of specialists and generalists. Access to complementary collaborators is an important and limited resource in academia, as in many other creative contexts. Past research has suggested that scientists choose collaborators strategically (Leahey and Reikowsky, 2008; Bikard, Murray, and Gans, 2015). Following our theoretical arguments, we expect specialists to be more sought after in faster-paced domains but generalists to be more sought after in slower-paced domains. We therefore expect an increase in collaboration levels of specialists (compared with generalists) in faster-paced domains of mathematics and a decline in their collaboration levels in slower-paced domains. Furthermore, we expect a decrease in collaboration of generalists in slower- versus faster-paced domains and a reverse effect for specialists.

[Insert table 5 about here]

To test these assertions, in table 5 we present results for the change in the collaboration rates of specialists and generalists after the fall of the Soviet Union. As before, we present results using our full sample in models 1 and 2 and results using our matched sample in models 3 and 4. As anticipated, the estimated β_2 suggests a relative decline of 7 percent and 10 percent in specialists' number of collaborators and specialists' number of unique collaborators, respectively, compared with generalists, in slower-paced domains. In faster-paced domains, however, specialists' total number of collaborators and the number of unique collaborators increased by up to 46 percent and 59 percent, respectively, compared with generalists, using their

differential change in the slower-paced domains as the baseline. Similarly, we observe declines of 31 and 33 percent in generalists' numbers of collaborators and numbers of unique collaborators, respectively, in faster- versus slower-paced domains, while specialists in faster-paced domains experience increases of 18 percent and 7 percent, respectively, over their specialist counterparts in slower-paced domains. Figures A2 and A3 in the Online Appendix show the timing of change in specialists' and generalists' collaboration rates. As before, we do not observe any collaboration trends in years before the Soviet collapse.

While these results show the change in collaboration rates of specialists and generalists after the Soviet collapse, they do not show changes in the compositions of collaborations—changes in the rates of specialist–specialist, specialist–generalist, and generalist–generalist collaborations. Unpacking the changes in collaboration composition is not empirically straightforward. Collaboration is a matching process that adds an additional layer of complexity to our estimations. For example, while one creative type (say, specialists) may decide to reduce its collaboration with the other (say, generalists) due to its lower benefits, the latter may put more effort in securing collaborations with the former due to its higher benefits. Hence it is difficult to make an ex-ante theoretical prediction about changes in some collaboration types. Moreover, because mathematicians on average have few collaborators, breaking down the small number of collaborators into different categories can lead to less accurate estimations with larger standard errors. Nonetheless, we provide some evidence of change in collaboration compositions in tables 6 and 7. The estimates suggest a statistically significant decline in collaboration among specialists and an increase in collaborations between specialists and generalists in slower-paced domains (β_2). In comparison, collaborations in the faster-paced environments do not experience such a change (β_1). In other words, specialists and generalists in faster-paced domains remain

more likely to collaborate with specialists and less likely to collaborate with generalists than their counterparts in slower-paced domains who prefer collaborating with generalists over specialists. Moreover, generalists in faster-paced environments reduce collaboration with other generalists when compared with generalists in slower-paced environments, while maintaining approximately the same levels of collaboration with specialists (β_3). At the same time, specialists in faster-paced domains increase collaboration with other specialists and decrease collaboration with generalists, when compared with specialists in slower-paced domains ($\beta_1 + \beta_3$), a statistically significant result. This finding is in line with our assertion that the observed creative advantages of specialists in faster-paced domains and of generalists in slower-paced domains are associated with a higher propensity of these individuals to be desired collaborators.

[Insert tables 6 and 7 about here]

We conducted several additional robustness checks to further corroborate our findings. One source of concern is that our results might be affected by an increase in labor market competition due to the increase in migration of Soviet mathematicians to other countries. To address this concern, we tested whether our results hold in geographical areas with little to no Soviet impact. Following the empirical strategy in Agrawal, Goldfarb, and Teodoridis (2016: page 113) we focused on Japan, “a country with no documented evidence of Soviet immigration in mathematics” and “which consistently ranks in the top ten mathematics research.” We find that all our results for generalists and specialists persist in the subset of Japanese-flagged authors, which indicates that our estimations are the result of the pace of the emergence of new knowledge components and not of labor market competition. The results are reported in tables A8 to A10 in the Online Appendix.

Moreover, we provide a battery of additional robustness checks for our estimations in the Online Appendix. Tables A11 to A13 show that our results hold if we use the 50-percent threshold on the diversification index to define specialists and generalists. In addition, we show robustness to the use of a continuous measure of specialization (tables A1 to A3) and to a continuous measure of the impact of the Soviet collapse on different domains of mathematics (tables A4 to A6). Finally, tables A14 to A16 show the sensitivity of our estimates to the exclusion of individual fixed effects and year fixed effects. Overall, the estimates and their interpretation are robust to these additional tests.

Discussion and Conclusion

Knowledge domains are not always stable. When the pace of change is slow, creative workers might need to reach beyond the traditional boundaries of their field to identify new recombination opportunities. When the pace of change accelerates, however, creative workers increasingly succeed by identifying and exploiting recombination opportunities emerging at the knowledge frontier. Both approaches can drive creative performance, but they constitute two different types of creative recombination that require very different skills. Individuals who spread their efforts across several domains and become generalists are likely to develop the capability to carry out recombinations across those domains efficiently. In contrast, individuals who concentrate their research efforts in a single knowledge domain and become specialists are likely to excel at taking advantage of new knowledge in their domain of specialty. The pace of change in a domain is therefore likely to affect the relative benefits of specialization in creative work.

Creative workers face a trade-off in deciding whether to focus their efforts on a narrow research domain or instead to spread their work across various fields. Both strategies present

advantages, and the superiority of each is a matter of debate in prior literature. We proposed that those disagreements might be driven in part by attempts to generalize from different domains exhibiting different paces of change. Just as the performance of firms' strategy is linked to the dynamics of the field in which the firms compete, the creative performance of individuals with different levels of specialization is likely to be linked to the dynamics of the knowledge domain in which they work. In particular, generalists are likely to perform better in slower-paced domains whereas specialists should perform best in faster-paced ones. Furthermore, those dynamics are likely to be amplified by within-domain competition for scarce resources. The performance of a specific creative strategy depends on the performance of competitors' strategy. Even though generalists in faster-paced domains are likely to have access to more recombination opportunities than generalists in slower-paced domains, they are likely to suffer much more from the competition of specialists.

We hypothesized and found empirical support that generalists perform relatively better than specialists in slower-paced environments but perform relatively worse as the pace of change increases. The Soviet collapse led to an unexpected and substantial acceleration of the pace of change in some subfields of theoretical mathematics but not in others. In the fields most affected, the performance of specialists improved sharply relative to that of generalists. At the same time, generalists performed relatively better than their specialist colleagues in the less affected fields, where the pace of change was slower. Differences in performance are visible across a variety of measures, including publication counts, citation-weighted publication counts, counts of breakthroughs, and even individual ability to attract collaborators. Furthermore, we find that the performance of generalists in affected fields decreases as the pace of change accelerates,

presumably because of the steep increase in the ability of specialist competitors to secure scarce resources in those domains.

One should note that our theoretical arguments do not rely on the type of sudden change in the knowledge frontier that we observe in our empirical setting. The collapse of the Soviet Union and the sudden influx of new knowledge in some areas of theoretical mathematics versus others is an essential part of our empirical strategy, but it is not required theoretically. This natural experiment helps us empirically isolate the variance in the pace of change across domains independent of the ex-ante activities of the creative workers in those domains. However, our theoretical claims apply to any other setting in which creative output relies on knowledge recombination.

Nevertheless, while the natural experiment of the unexpected Soviet collapse provides a rewarding test for our theoretical predictions, some limitations remain. First, despite the richness of our data and the comprehensive role of theoretical mathematics in creative work across a multitude of areas, we studied one setting. The trade-offs associated with being specialists or generalists might be different in other creative settings. Moreover, we focus here on a specific type of change. At times, some discoveries challenge the very foundations of entire knowledge domains, provoking what Kuhn (1970) referred to as scientific revolutions. For example, the introduction of Einsteinian dynamics in theoretical physics challenged many of the assumptions held by Newtonian physicists. Past research suggested that specialists in a domain that has experienced a scientific revolution are more likely to resist adapting to the foundational changes in their domain of expertise (Kuhn, 1970: 151). In cases of scientific revolution in a domain, specialists' domain-specific heuristics and problem-solving skills may no longer give them any advantage over generalists, as they may all be challenged by the radical changes in the

foundations of the field. Thus we make no strong claims of generalizability and hope that future research will explore whether different types of change might have different consequences for the performance of specialists and generalists.

Second, our study is limited by our somewhat static operationalization of the generalist and specialist creative strategies. The ability of individuals to become generalists or specialists varies. Moreover, the distinction between specialists is not always as clear as it appears in theoretical mathematics. The implications of our results for individuals specializing in tools or topics that have broad applications (e.g., general purpose technologies) remain unclear. Creative workers might also become specialists along one dimension and generalists along another one (Kacperczyk and Younkin, 2017). Besides, individuals might shift strategy over the span of their career (Mannucci and Yong, 2017). For example, junior researchers might exhibit greater specialization whereas senior individuals might exhibit greater diversification. In our empirical analysis, we controlled for a quadratic effect of age to account for this possibility. But this approach does not consider that junior specialists might become increasingly diverse as they advance in their careers. At the same time, it is unclear whether generalists might narrow their focus as they encounter a prolific area of research. To address this concern, we calculated our index of diversification on a rolling basis to seek evidence of significant changes in diversification at the individual level throughout the course of our dataset. We did not find such evidence.

Our study makes several theoretical contributions. First, we describe some of the trade-offs of specialization in creative work. Before this study, a large stream of work described the advantages of being a generalist, highlighting the benefits of brokering otherwise distant knowledge components (Hargadon and Sutton, 1997; Uzzi and Spiro, 2005; Audia and Goncalo,

2007; Cattani and Ferriani, 2008; Jeppesen and Lakhani, 2010). Other studies instead emphasized benefits of specialization such as specialists' deeper understanding of their knowledge domain and their clearer identity (Birnbaum, 1981; Leahey, 2007; Jones, 2009; Conti, Gambardella, and Mariani, 2013). We extend this literature by highlighting that the decision of whether to become a specialist or a generalist is in fact a strategic one and that the creative performance of generalists and specialists depends on the context in which they operate.

Second, our findings highlights the existence of two types of recombination: "horizontal recombination" in which creative workers broker components across knowledge domains and "vertical recombination" in which they take advantage of the forward movement of the knowledge frontier. This distinction provides a new perspective on the literature on distant versus local search (Fleming, 2001; Kaplan and Vakili, 2015; Leahey, Beckman, and Stanko, 2017). Local recombinations are often described as less risky and somehow easier than more distant ones. The distant versus local terminology might therefore lead to overstating the value of the work of generalists while understating that of specialists. Instead, our study highlights that recombination based on newly emerging knowledge in a faster-paced domain is not as trivial as prior literature suggests. It involves important skills to identify those opportunities early on and to exploit them efficiently. The role of those skills for creative success has received little attention to date, but our study highlights that specialists tend to outperform their generalist peers in this respect.

Third, our study highlights the crucial role of competitive dynamics in creative work. Prior research described such a dynamic in the case of individual social status in science. The typical argument is that high-status individuals can reap more rewards from the products of their work, which further increases their performance at the expense of their competitors' (e.g.,

Reschke, Azoulay, and Stuart, 2017). Our study extends this stream of research by highlighting that the same competitive dynamic can be triggered by variations in the pace of change in a knowledge domain. More specifically, we find that the same creative strategies can lead to different levels of performance depending on the success of the competitors' strategy. Studies of creative performance that overlook these competitive dynamics might therefore lead to erroneous conclusions.

Fourth, and related to the issue of competition, our study contributes to the literature on collaboration in knowledge creation by contextualizing the common finding that collaboration is associated with high creative performance. Past research has emphasized the key role of collaboration in fostering creativity by facilitating more diverse knowledge recombinations and more efficient selection of good ideas (Reagans, Zuckerman, and McEvily, 2004; Fleming, Mingo, and Chen, 2007; Singh and Fleming, 2010). Yet these studies usually overlook that collaboration is often a choice. Our findings therefore contribute to a growing literature exploring the determinants of collaboration strategies in creative work (Leahey and Reikowsky, 2008; Bikard, Murray, and Gans, 2015; Bikard, Vakili, and Teodoridis, 2018) by highlighting how the relative cognitive advantage of individuals in a domain shapes their collaboration opportunities. We also highlight how preferential access to complementary collaborators can reinforce the creative advantage. More broadly, our results call for more research to understand the dynamics of competition for collaboration and the complex market for collaborators.

Additionally, our study has implications for the organization of firms' R&D. The contribution of generalists and specialists to R&D performance is distinct but potentially complementary. While generalists might draw on components located beyond the traditional knowledge domain of the firm, specialists are likely to be better able to take advantage of the

emergence of new components. Both sets of skills might be valuable. In fact, the differences between the two creative strategies can provide opportunities for productive collaborations between specialists and generalists (Teodoridis, 2017). Our study suggests that the appropriate balance of specialists and generalists inside the firm will depend not only on its intention to absorb external knowledge from within its area of specialization or beyond it but also on the pace of change in those knowledge domains.

Our study constitutes a first step in highlighting how the pace of change shapes the performance of creative workers. It also highlights important trade-offs associated with specialization in creative work. The importance of furthering this line of research should not be understated. Creativity and innovation play a growing role in individual and firm performance, and there is no sign that the pace of economic change might stop evolving differently across domains and over time. By highlighting the fact that creative workers rarely evolve in static knowledge domains, we hope our study enhances our understanding of the drivers of creative performance and triggers future research on creative strategies in a changing world.

Acknowledgments

We would like to thank seminar participants at London Business School, Rotman School of Management at University of Toronto, University of Bath School of Management, ESMT Berlin, Grenoble School of Management, Strategy Research Forum, Wharton Strategy Science Conference, and Academy of Management Conference for helpful comments and suggestions. We would also like to thank senior editor Christine Beckman and the three anonymous reviewers for their valuable and constructive feedback on the paper. This study was supported by funding from the London Business School.

Supplemental Material

Supplemental material for this article can be found in the Online Appendix at <http://journals.sagepub.com/doi/suppl/10.1177/0001839218793384>.

REFERENCES

- Aarts, H., and A. Dijksterhuis
2000 "Habits as knowledge structures: Automaticity in goal-directed behavior." *Journal of Personality and Social Psychology*, 78: 53–63.
- Aarts, H., B. Verplanken, and A. van Knippenberg
1998 "Predicting behavior from actions in the past: Repeated decision making or a matter of habit?" *Journal of Applied Social Psychology*, 28: 1355–1374.
- Agrawal, A., A. Goldfarb, and F. Teodoridis
2016 "Understanding the changing structure of scientific inquiry." *American Economic Journal: Applied Economics*, 8: 100–128.
- Ahuja, G., and C. M. Lampert
2001 "Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions." *Strategic Management Journal*, 22: 521–543.
- Amabile, T. M.
1983 "The social psychology of creativity: A componential conceptualization." *Journal of Personality and Social Psychology*, 45: 357–377.
- Audia, P. G., and J. A. Goncalo
2007 "Past success and creativity over time: A study of inventors in the hard disk drive industry." *Management Science*, 53: 1–15.
- Azoulay, P., T. Stuart, and Y. Wang
2013 "Matthew: Effect or fable?" *Management Science*, 60: 92–109.
- Barnett, W. P.
1990 "The organizational ecology of a technological system." *Administrative Science Quarterly*, 35: 31–60.
- Bikard, M., F. Murray, and J. S. Gans
2015 "Exploring trade-offs in the organization of scientific work: Collaboration and scientific reward." *Management Science*, 61: 1473–1495.
- Bikard, M., K. Vakili, and F. Teodoridis
2018 "When collaboration bridges institutions: The impact of university–industry collaboration on academic productivity." *Organization Science* (in press).
- Bilalić, M., P. McLeod, and F. Gobet

2008 “Inflexibility of experts—Reality or myth? Quantifying the Einstellung Effect in chess masters.” *Cognitive Psychology*, 56: 73–102.

Birnbaum, P. H.

1981 “Integration and specialization in academic research.” *Academy of Management Journal*, 24: 487-503.

Blackwell, M., S. Iacus, G. King, and G. Porro

2009 “Cem: Coarsened exact matching in Stata.” *Stata Journal*, 9: 524–546.

Boring, E. G.

1955 “Dual role of the zeitgeist in scientific creativity.” *Scientific Monthly*, 80: 101–106.

Borjas, G. J., and K. B. Doran

2012 “The collapse of the Soviet Union and the productivity of American mathematicians.” *Quarterly Journal of Economics*, 127: 1143-1203.

Burt, R. S.

2004 “Structural holes and good ideas.” *American Journal of Sociology*, 110: 349–399.

Carnabuci, G., and J. Bruggeman

2009 “Knowledge specialization, knowledge brokerage and the uneven growth of technology domains.” *Social Forces*, 88: 607–641.

Cattani, G., and S. Ferriani

2008 “A core/periphery perspective on individual creative performance: Social networks and cinematic achievements in the Hollywood film industry.” *Organization Science*, 19: 824–844.

Chai, S.

2017 “Near misses in the breakthrough discovery process.” *Organization Science*, 28: 411–428.

Chase, W. G., and H. A. Simon

1973 “Perception in chess.” *Cognitive Psychology*, 4: 55–81.

Chi, M. T., R. Glaser, E. Rees, and R. Sternberg

1982 “Expertise in problem solving.” In R. Sternberg (ed.), *Advances in the Psychology of Human Intelligence*, [vol. 1, pp. 7-76]. Hillsdale, NJ: Lawrence Erlbaum Associates.

Conti, R., A. Gambardella, and M. Mariani

2013 “Learning to be Edison: Inventors, organizations, and breakthrough inventions.” *Organization Science*, 25: 833–849.

Dane, E.

2010 “Reconsidering the trade-off between expertise and flexibility: A cognitive Entrenchment perspective.” *Academy of Management Review*, 35: 579–603.

Dosi, G.

1982 “Technological paradigms and technological trajectories : A suggested interpretation of the determinants and directions of technical change.” *Research Policy* 11: 147–162.

Dougherty, D.

1992 “Interpretive barriers to successful product innovation in large firms.” *Organization Science* 3: 179–202.

Dunbar, K.

1995 “How scientists really reason: Scientific reasoning in real-world laboratories.” *The Nature of Insight*, 18: 365–395.

Ferguson, J.-P., and G. Carnabuci

2017 “Risky recombinations: Institutional gatekeeping in the innovation process.” *Organization Science*, 28: 133–151.

Fleming, L.

2001 “Recombinant uncertainty in technological search.” *Management Science* 47: 117–132.

Fleming, L.

2002 “Finding the organizational sources of technological breakthroughs: The story of Hewlett-Packard’s thermal ink-jet.” *Industrial and Corporate Change* 11: 1059–1084.

Fleming, L., S. Mingo, and D. Chen

2007 “Collaborative brokerage, generative creativity, and creative success.” *Administrative Science Quarterly*, 52: 443–475.

Franzoni, C., G. Scellato, and P. Stephan

2011 “Changing incentives to publish.” *Science*, 333: 702–703.

Frensch, P. A., and R. J. Sternberg

1989 “Expertise and intelligent thinking: When is it worse to know better.” In R. Sternberg (ed.), *Advances in the Psychology of Human Intelligence*, 5: 157–188. Hillsdale, NJ: Erlbaum.

Furman, J., and S. Stern

2011 “Climbing atop the shoulders of giants: The impact of institutions on cumulative research.” *American Economic Review*, 101: 1933–1963.

Ganguli, I.

2014 “Scientific brain drain and human capital formation after the end of the Soviet Union.” *International Migration*, 52: 95–110.

- Gould, S. J., and N. Eldredge
1977 "Punctuated equilibria: The tempo and mode of evolution reconsidered." *Paleobiology*, 3: 115–151.
- Graham, L. R.
1993 *Science in Russia and the Soviet Union: A Short History*. Cambridge: Cambridge University Press.
- Graham, L. R., and I. Dezhina
2008 *Science in the New Russia: Crisis, Aid, Reform*. Bloomington, IN: Indiana University Press.
- Hargadon, A., and Y. Douglas
2001 "When innovations meet institutions: Edison and the design of the electric light." *Administrative Science Quarterly*, 46: 476–501.
- Hargadon, A., and R. Sutton
1997 "Technology brokering and innovation in a product development firm." *Administrative Science Quarterly* 42: 716–749.
- Hauptman, O., and K. K. Hirji
1996 "The influence of process concurrency on project outcomes in product development: An empirical study of cross-functional teams." *IEEE Transactions on Engineering Management*, 43: 153–164.
- Henderson, R., and K. Clark
1990 "Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms." *Administrative Science Quarterly*, 35: 9–30.
- Iacus, S. M. and King G. Porro
2011 "Multivariate matching methods that are monotonic imbalance bounding." *Journal of American Statistics Association*, 106(493): 345-361.
- Jeppesen, L. B., and K. R. Lakhani
2010 "Marginality and problem-solving effectiveness in broadcast search." *Organization Science*, 21: 1016–1033.
- Jones, B.
2009 "The burden of knowledge and the 'death of the Renaissance man': Is innovation getting harder?" *Review of Economic Studies*, 76: 283–317.
- Jones, B., and B. Weinberg
2011 "Age dynamics in scientific creativity." *Proceedings of the National Academy of Sciences*, 108: 18910–18914.
- Kacperczyk, A., and P. Younkin

2017 “The paradox of breadth: The tension between experience and legitimacy in the transition to entrepreneurship.” *Administrative Science Quarterly*, 62: 731–764.

Kaplan, S., and K. Vakili

2015 “The double-edged sword of recombination in breakthrough innovation.” *Strategic Management Journal*, 36: 1435–1457.

Knorr-Cetina, K.

1999 *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press.

Knudsen, T., and K. Srikanth

2014 “Coordinated exploration: Organizing joint search by multiple specialists to overcome mutual confusion and joint myopia.” *Administrative Science Quarterly*, 59: 409–441.

Kuhn, T. S.

1970 *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Latour, B., and S. Woolgar

1986 *Laboratory Life: The Construction of Scientific Facts*. Princeton, NJ: Princeton University Press.

Leahey, E.

2007 “Not by productivity alone: How visibility and specialization contribute to academic earnings.” *American Sociological Review*, 72: 533–561.

Leahey, E., C. M. Beckman, and T. L. Stanko

2017 “Prominent but less productive: The impact of interdisciplinarity on scientists’ research.” *Administrative Science Quarterly*, 62: 105–139.

Leahey, E., and R. C. Reikowsky

2008 “Research specialization and collaboration patterns in sociology.” *Social Studies of Science*, 38:425-440.

Levinthal, D. A.

1998 “The slow pace of rapid technological change: Gradualism and punctuation in technological change.” *Industrial and Corporate Change*, 7: 217–247.

Lim, K.

2009 “The many faces of absorptive capacity: Spillovers of copper interconnect technology for semiconductor chips.” *Industrial and Corporate Change*, 18: 1249-1284.

Luchins, A. S.

1942 “Mechanization in problem solving: The effect of Einstellung.” *Psychological Monographs*, 54: i-95.

- Mannucci, P. V., and K. Yong
2017 “The differential impact of knowledge depth and knowledge breadth on creativity over individual careers.” *Academy of Management Journal* (in press). Published online ahead of print: <https://doi.org/10.5465/amj.2016.0529>.
- Merton, R. K.
1957 “Priorities in scientific discovery: A chapter in the sociology of science.” *American Sociological Review*, 22: 635–659.
- Merton, R. K.
1968 “The Matthew Effect in science The reward and communication systems of science are considered.” *Science*, 159: 56-63.
- Murray, K. B., and G. Häubl
2007 “Explaining cognitive lock-in: The role of skill-based habits of use in consumer choice.” *Journal of Consumer Research*, 34: 77–88.
- Nagle, F., and F. Teodoridis
2017 “Jack of all trades and master of knowledge: The role of generalists in novel knowledge integration” USC-INET Research Paper No. 17-23.
- Orlikowski, W. J.
2002 “Knowing in practice: Enacting a collective capability in distributed organizing.” *Organization Science*, 13: 249–273.
- Phene, A., K. Fladmoe-Lindquist, and L. Marsh
2006 “Breakthrough innovations in the U.S. biotechnology industry: The effects of technological space and geographic origin.” *Strategic Management Journal*, 27: 369–388.
- Polanyi, M.
1958 *Personal Knowledge*. London: Routledge.
- Reagans, R., E. Zuckerman, and B. McEvily
2004 “How to make the team: Social networks vs. demography as criteria for designing effective teams.” *Administrative Science Quarterly*, 49: 101–133.
- Reagans, R., and E. W. Zuckerman
2001 “Networks, diversity, and productivity: The social capital of corporate R&D teams.” *Organization Science*, 12: 502–517.
- Reschke, B. P., P. Azoulay, and T. E. Stuart
2017 “Status spillovers: The effect of status-conferring prizes on the allocation of attention.” *Administrative Science Quarterly* (in press). Published online ahead of print: <https://doi.dox.org/10.1177/0001839217731997>.
- Romer, P. M.

- 1990 “Endogenous technological change.” *Journal of Political Economy*, 98: S71-S102.
- Russo E.
2003 “Special report: The birth of biotechnology.” *Nature*, 421:456–457.
- Schilling, M. A.
2005 “A ‘small-world’ network model of cognitive insight.” *Creativity Research Journal*, 17: 131–154.
- Schilling, M. A., and E. Green
2011 “Recombinant search and breakthrough idea generation: An analysis of high impact papers in the social sciences.” *Research Policy*, 40: 1321–1331.
- Schumpeter, J. A.
1942 *Capitalism, Socialism and Democracy*. New York: Harper & Row.
- Silva, J. M. C. S., and S. Tenreiro
2006 “The log of gravity.” *Review of Economics and Statistics*, 88: 641–658.
- Singh, J., and L. Fleming
2010 “Lone inventors as sources of breakthroughs: Myth or reality?” *Management Science*, 56: 41–56.
- Song, X. M., and M. M. Montoya-Weiss
1998 “Critical development Activities for really new versus incremental products.” *Journal of Product Innovation Management*, 15: 124–135.
- Song, X. M., R. J. Thieme, and J. Xie
1998 “The impact of cross-functional joint involvement across product development stages: An exploratory study.” *Journal of Product Innovation Management*, 15: 289–303.
- Stephan, P.
2012 *How Economics Shapes Science*. Cambridge, MA: Harvard University Press.
- Taylor, A., and H. R. Greve
2006 “Superman or the Fantastic Four? Knowledge combination and experience in innovative teams.” *Academy of Management Journal*, 49: 723-740.
- Teodoridis, F.
2017 “Understanding team knowledge production: The interrelated roles of technology and expertise.” *Management Science* (in press). Published online ahead of print: <https://doi.org/10.1287/mnsc.2017.2789>.
- Tikhomirov, V. M.

2007 “On Moscow mathematics—Then and now.” In Zdravkovska, S., Duren, P. L. (eds), *Golden Years of Moscow Mathematics, 273-283*. Providence, RI: American Mathematical Society.

Tushman, M. L., and P. Anderson

1986 “Technological discontinuities and organizational environments.” *Administrative Science Quarterly*, 31: 439–465.

Tyre, M. J., and O. Hauptman

1992 “Effectiveness of organizational responses to technological change in the production process.” *Organization Science*, 3: 301–320.

Uzzi, B., S. Mukherjee, M. Stringer, and B. Jones

2013 “Atypical combinations and scientific impact.” *Science*, 342: 468–472.

Uzzi, B., and J. Spiro

2005 “Collaboration and creativity: The small world Problem.” *American Journal of Sociology*, 111: 447–504.

Vakili, K., and A. M. McGahan

2016 “Health care’s grand challenge: Stimulating basic science on diseases that primarily afflict the poor.” *Academy of Management Journal*, 59: 1917–1939.

Weisberg, R. W.

1998 “Creativity and knowledge: A challenge to theories.” In R. J. Sternberg (ed.), *Handbook of Creativity*, 226-250. Cambridge: Cambridge University Press.

Weitzman, M. L.

1998 “Recombinant growth.” *Quarterly Journal of Economics*, 113: 331–360.

Authors' Biographies

Florenta Teodoridis is an Assistant Professor of Strategy at the Marshall School of Business, University of Southern California. Florenta's research interests are in the economics of innovation. Her research agenda includes studies analyzing the interplay between innovators and technology as input and output of knowledge production, factor influencing creativity and collaboration, and the relationship between knowledge advancement and markets. USC Marshall School of Business, 701 Exposition Blvd., Los Angeles, CA 90089, USA (e-mail: florenta.teodoridis@marshall.usc.edu).

Michaël Bikard is an Assistant Professor of Strategy and Entrepreneurship at London Business School. He received his Ph.D. from MIT Sloan School of Management. His research interests include creativity, innovation, and technology strategy. London Business School, Sussex Place, London, NW1 4SA, United Kingdom (e-mail: mbikard@london.edu).

Keyvan Vakili is an Assistant Professor of Strategy and Entrepreneurship at London Business School. He received his PhD from Rotman School of Management, University of Toronto. His research interests include creativity, knowledge recombination, innovation and technology strategy, and the relationship between innovation and collaboration. London Business School, Sussex Place, London, NW1 4SA, United Kingdom (e-mail: kvakili@london.edu).

Table 1. Subfield Rank of Soviet Contributions to Theoretical Mathematics*

Subfield rank	Theoretical mathematics category	Description
1	Analysis	Integral equations
2	Analysis	Fourier analysis
3	Analysis	Partial differential equations
4	Analysis	Sequences, series, summability
5	Analysis	Potential theory
6	Analysis	Calculus of variations and optimal control; optimization
7	Analysis	Integral transforms, operational calculus
8	Analysis	Functions of a complex variable
9	Algebra	General algebraic systems
10	Analysis	Difference equations and functional equations
11	Analysis	Operator theory
12	Algebra	Non-associative rings and non-associative algebras
13	Analysis	Approximations and expansions
14	Geometry	Global analysis, analysis on manifolds
15	Analysis	Several complex variables and analytic spaces
16	Analysis	Special functions
17	Algebra	Topological groups, lie groups, and analysis upon them
18	Geometry	General topology
19	Algebra	Group theory and generalizations
20	Algebra	Measure and integration
21	Algebra	Category theory; homological algebra
22	Analysis	Algebraic topology
23	Algebra	Real functions, including derivatives and integrals
24	Geometry	Convex geometry and discrete geometry
25	Algebra	Algebraic geometry
26	Analysis	Abstract harmonic analysis
27	Algebra	Linear and multilinear algebra; matrix theory
28	Algebra	Order theory
29	Algebra	Field theory and polynomials
30	Algebra	Combinatorics
31	Geometry	Geometry
32	Geometry	Manifolds
33	Algebra	Commutative rings and algebras

* The ranking on the left indicates the level of impact of the fall of the Soviet Union on the subfield. The higher the subfield's ranking, the more it was affected by the shock. The ranking is based on Agrawal, Goldfarb, and Teodoridis (2016).

Table 2. Summary Statistics for the Full Sample and the Matched Sample (1980–2000)

Variable	Mean	S.D.
Panel A: Full Sample (N = 123,139)		
Citation-weighted number of publications per year	5.724	35.193
Simple publication count per year	.732	1.299
Number of breakthrough publications (in top 5% cited) per year	.043	.275
Number of breakthrough publications (in top 10% cited) per year	.090	.402
Number of collaborators between 1980 and 1988	.541	1.519
Number of unique collaborators between 1980 and 1988	.413	.937
Panel B: Matched Sample (N = 81,762)		
Citation-weighted number of publications per year	4.678	20.178
Simple publication count per year	.690	1.151
Number of breakthrough publications (in top 5% cited) per year	.037	.229
Number of breakthrough publications (in top 10% cited) per year	.081	.346
Number of collaborators between 1980 and 1988	.505	1.327
Number of unique collaborators between 1980 and 1988	.412	.920

Table 3. Specialists versus Generalists before the Collapse of the Soviet Union (1980–1989)*

Variable	Specialists	Generalists	t-test difference
Panel A: Full Sample			
Number of mathematicians	4,042	2,213	
Total citation-weighted number of publications between 1980 and 1988	44.777 (108.712)	61.568 (148.799)	-16.791**
Total number of publications between 1980 and 1988	8.047 (9.413)	8.856 (10.386)	-.809*
Total number of breakthrough publications (in top 5% cited) between 1980 and 1988	.356 (1.116)	.560 (1.579)	-.204**
Total number of breakthrough publications (in top 10% cited) between 1980 and 1988	.780 (1.710)	1.102 (2.244)	-.322**
Average number of collaborators between 1980 and 1988	.703 (1.386)	.588 (.729)	.115**
Average number of unique collaborators between 1980 and 1988	.446 (.572)	.480 (.517)	-.034*
Panel B: Matched Sample			
Number of mathematicians	2,038	2,038	
Total citation-weighted number of publications between 1980 and 1988	39.915 (66.830)	41.728 (67.621)	-1.813
Total number of publications between 1980 and 1988	7.261 (5.395)	7.337 (5.323)	-.075
Total number of breakthrough publications (in top 5% cited) between 1980 and 1988	.340 (.910)	.376 (.891)	-.036
Total number of breakthrough publications (in top 10% cited) between 1980 and 1988	.816 (1.537)	.835 (1.490)	-.019
Average number of collaborators between 1980 and 1988	.528 (.772)	.533 (.631)	.018
Average number of unique collaborators between 1980 and 1988	.410 (.520)	.448 (.478)	.037*

* $p < .05$; ** $p < .01$.

* Standard deviations are in parenthesis in Column 1 and 2, and p-values are in parentheses in Column 3.

Table 4. Changes in the Publication Output of Specialist and Generalist Mathematicians after the Collapse of the Soviet Union*

Variable	Full Sample		Matched Sample	
	Simple count of publications	Citation-weighted count of publications	Simple count of publications	Citation-weighted count of publications
	(1)	(2)	(3)	(4)
Specialist \times SovietImpact \times AfterSovietCollapse (β_1)	.315** (.113)	.605** (.205)	.390** (.129)	.777** (.253)
Specialist \times AfterSovietCollapse (β_2)	-.078* (.032)	-.254* (.103)	-.056 (.038)	-.241** (.080)
SovietImpact \times AfterSovietCollapse (β_3)	-.264** (.099)	-.454** (.176)	-.382** (.114)	-.610** (.208)
No. of observations	113,512	113,406	76,795	76,783
No. of mathematicians	6,140	6,132	4,024	4,024
Chi ²	1059.76**	203.81**	663.53**	104.60**
Log-likelihood	-104275.08	-786985.39	-68821.16	-451570.52

* $p < .05$; ** $p < .01$.

* The data are a panel at the author level based on publication data from 1980 through 2000. The unit of analysis is the author-year. All models are conditional fixed-effect Poisson with robust standard errors, clustered at the author level, in parentheses. All models include controls for cumulative publications, nonlinear age profile, and individual and year fixed effects. The difference in the number of observations across models is a consequence of estimating all our models using the *xtpoisson* command in Stata; the command drops units without within-individual variance after factoring in all the independent and control variables.

Table 5. Changes in the Collaboration Rates of Specialist and Generalist Mathematicians after the Collapse of the Soviet Union*

Variable	Full Sample		Matched Sample	
	Total number of collaborators (1)	Total number of unique collaborators (2)	Total number of collaborators (3)	Total number of unique collaborators (4)
Specialist × SovietImpact × AfterSovietCollapse (β_1)	.428** (.157)	.340* (.139)	.457* (.179)	.380* (.157)
Specialist × AfterSovietCollapse (β_2)	-.110* (.046)	-.074* (.037)	-.100+ (.054)	-.057 (.044)
SovietImpact × AfterSovietCollapse (β_3)	-.260* (.132)	-.270* (.119)	-.396* (0.155)	-.371** (.135)
No. of observations	96,917	96,917	65,986	65,986
No. of mathematicians	5,243	5,243	3,459	3,459
Chi ²	122.91**	196.63**	63.56**	116.44**
Log-likelihood	88857.85	-71016.25	-57694.92	-48285.65

+ $p < .10$; * $p < .05$; ** $p < .01$.

* The data are a panel at the author level based on publication data from 1980 through 2000. The unit of analysis is the author-year. All models are conditional fixed-effect Poisson with robust standard errors, clustered at the author level, in parentheses. All models include controls for cumulative publications, nonlinear age profile, and individual and year fixed effects. The difference in the number of observations across models is a consequence of estimating all our models using the *xtpoisson* command in Stata; the command drops units without within-individual variance after factoring in all the independent and control variables.

Table 6. Changes in Collaboration with Specialist Mathematicians after the Collapse of the Soviet Union*

Variable	Full Sample		Matched Sample	
	Total number of collaborators (1)	Total number of unique collaborators (2)	Total number of collaborators (3)	Total number of unique collaborators (4)
Specialist × SovietImpact × AfterSovietCollapse (β_1)	.195 (.412)	.348 (.356)	.191 (.484)	.399 (.397)
Specialist × AfterSovietCollapse (β_2)	-.621** (.092)	-.541** (.071)	-.616** (.101)	-.526** (.081)
SovietImpact × AfterSovietCollapse (β_3)	.071 (.397)	-.213 (.345)	-.013 (.466)	-.299 (.381)
No. of observations	53,812	53,812	34,459	34,459
No. of mathematicians	2,905	2,905	1,798	1,798
Chi ²	372.09**	540.51**	203.90**	266.56**
Log-likelihood	-28263.91	-21838.02	-16367.74	-13485.23

* $p < .05$; ** $p < .01$.

* The data are a panel at the author level based on publication data from 1980 through 2000. The unit of analysis is the author-year. All models are conditional fixed-effect Poisson with robust standard errors, clustered at the author level, in parentheses. All models include controls for cumulative publications, nonlinear age profile, and individual and year fixed effects. The difference in the number of observations across models is a consequence of estimating all our models using the *xtpoisson* command in Stata; the command drops units without within-individual variance after factoring in all the independent and control variables.

Table 7. Changes in Collaboration with Generalist Mathematicians after the Collapse of the Soviet Union*

Variable	Full Sample		Matched Sample	
	Total number of collaborators (1)	Total number of unique collaborators (2)	Total number of collaborators (3)	Total number of unique collaborators (4)
Specialist × SovietImpact × AfterSovietCollapse (β_1)	.177 (.263)	.117 (.247)	.758* (.317)	.521+ (.306)
Specialist × AfterSovietCollapse (β_2)	.416** (.088)	.402** (.075)	.276* (.115)	.340** (.100)
SovietImpact × AfterSovietCollapse (β_3)	-.530** (.185)	-.444* (.172)	-.797** (.200)	-.658** (.192)
No. of observations	44,207	44,207	33,611	33,611
No. of mathematicians	2,352	2,352	1,754	1,754
Chi ²	274.43**	332.42**	251.86**	287.85**
Log-likelihood	-	-	-14616.25	-12246.95
	19700.48	15988.21		

+ $p < .10$; * $p < .05$; ** $p < .01$.

* The data are a panel at the author level based on publication data from 1980 through 2000. The unit of analysis is the author-year. All models are conditional fixed-effect Poisson with robust standard errors, clustered at the author level, in parentheses. All models include controls for cumulative publications, nonlinear age profile, and individual and year fixed effects. The difference in the number of observations across models is a consequence of estimating all our models using the *xtpoisson* command in Stata; the command drops units without within-individual variance after factoring in all the independent and control variables.

Figure 1. Adjusted average citation-weighted number of publications by specialist and generalist mathematicians in faster-paced and slower-paced areas after the collapse of the Soviet Union.





Figure 2. Estimated relative difference in the citation-weighted number of publications of specialists versus generalists after the collapse of the Soviet Union.*



* We base this figure on 10 years of publication data before the collapse of the Soviet Union and 10 years after the collapse. Each point on graph (a) represents the coefficient value on the covariate *Specialist × TimePeriod* and thus describes the relative difference in quality-adjusted publication rates between specialists and generalists in slower-paced areas. Each point on graph (b) represents the coefficient value on the covariate *Specialist × SovietImpact × TimePeriod* and thus describes the relative difference in quality-adjusted publication rates between specialists and generalists in faster-paced areas and the same difference in slower-paced areas. Each point on graph (c) represents the coefficient value on the covariate *SovietImpact × TimePeriod* and thus describes the relative difference in quality-adjusted publication rates between generalists in faster- versus slower-paced areas. Each point on graph (d) represents the sum of coefficients $\beta_1 + \beta_3$ and thus describes the relative difference in quality-adjusted publication rates between specialists in faster- versus slower-paced areas. The bars surrounding each point represent the 95% confidence interval. Note that the larger confidence intervals are due to reduced degrees of freedom, as we split the post-Soviet dummy into multiple period dummies. All values are relative to the base-year group of 1987–1989.