Can Specialization Foster Creativity?
Mathematics and the Collapse of the Soviet Union

Florenta Teodoridis
USC Marshall School of Business
florenta.teodoridis@marshall.usc.edu

Keyvan Vakili
London Business School
kvakili@london.edu

Michaël Bikard
London Business School
mbikard@london.edu

January 2017

Preliminary: Please do not cite or circulate

Abstract

Past research emphasizes that high levels of specialization can decrease creative performance because creative workers end up reusing familiar components in familiar ways. We argue that this view of specialization might be overly pessimistic. At times, knowledge domains evolve fast – that is, new components become rapidly available. In such situations, deep expertise places specialists in an advantageous position to identify emerging creative opportunities. To test our theory, we exploit the collapse of the Soviet Union as a natural experiment that suddenly moved the frontier in some domains of theoretical mathematics more than in others. We find that specialists in domains that advanced the most produced not only more publications but also more breakthroughs than their generalist colleagues did. Specialists also increased their rate of collaboration. Thus, specialization seems to be a particularly advantageous approach to creative work in fast-evolving knowledge domains.

Despite the widely recognized benefits of specialization in economic life, its value in creative work is contentious. On one hand, creativity requires some domain-specific knowledge (Amabile, 1983; Mokyr, 2002; von Hippel, 1988). On the other hand, specialization can blind creative workers to unorthodox approaches. The risk is that specialists end up reusing the same set of familiar components in the same, familiar ways. Eventually, they run out of impactful recombinations, producing only marginal innovations (Fleming, 2001). To avoid the drawbacks of specialization, a large literature suggests, creative workers should seek access to diverse knowledge bases. Individuals who adopt this strategy – generalists, for brevity's sake[1] – arguably have more opportunities for novel knowledge recombinations as well as a wider set of heuristics that can help them break away from traditional thought patterns (Hargadon and Sutton, 1997; Benner and Tushman, 2002; Burt, 2004; Uzzi and Spiro, 2005; Fleming, Mingo, and Chen 2007; Taylor and Greve, 2006; Cattani and Ferriani, 2008).

An important implicit assumption underlying this line of thought is the relative stability of domain-specific knowledge components. Once most recombinations within a knowledge domain are exhausted, fresh recombinations and breakthrough innovations can presumably only occur through the importing of knowledge from distant domains. However, knowledge domains are rarely static. Some knowledge domains advance slowly at times, but many undergo periods of rapid, substantial advancements due to scientific and technological discoveries or to spontaneous surfacing of previously inaccessible knowledge. For example, Lim (2009) documents 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). Similar movements in knowledge frontiers have been documented in industries such as pharmaceuticals (Gambardella, 1992), materials (Klincewicz, 2016), and computing (Marx, Gans, and Hsu, 2014; Forman, Goldfarb, and Greenstein, 2016). Similarly, historical events have opened

[1] In this paper, we define specialists as individuals who have experience and expertise in a narrowly defined domain of knowledge. In contrast, a comparable generalist is an individual with the same amount of experience spread across multiple related or unrelated knowledge domains. The notion of specialist in our paper is very close to the notion of expert in some other studies. In principle, a specialist is an individual who has expertise in a domain. A generalist with the same amount of experience spread across multiple domains would, by definition, have less expertise in each domain. We provide examples of specialists and generalists in our setting in future sections.

the door to previously inaccessible knowledge. As an example, the end of World War II led to the release of various important, and previously secret, scientific advancements both by Nazi researchers and by the Allies to the rest of the world, inducing the rapid evolution of a number of knowledge domains including radar technology, medical technology, psychology, ballistic missiles, and cryptography (Hactcup, 2000; Ifrah, 2000; Samelson, 1977).

In this paper, we analyze the value of specialization for creative work in fast-evolving knowledge domains.[2] We argue that when the knowledge in a domain expands, specialists are more likely to achieve higher levels of output and to generate more impactful innovations. Specialization allows one to have a more thorough understanding of the knowledge landscape and a better grasp of the fundamental debates and gaps in the domain. Specialization also improves domain-specific problem-solving skills. In contrast, generalists' ability to span various knowledge domains is likely to mean that they are less embedded within any given domain. Therefore, while generalists may be more capable of bridging multiple knowledge areas and of producing impactful creative output when knowledge domains are relatively stable, they may be at a disadvantage for identifying, absorbing, and exploiting new knowledge when a knowledge domain quickly expands. Specialists, on the other hand, are not only in a better position to identify, understand, and exploit the new knowledge at the frontier of their respective domain of specialty, but are also better equipped to rapidly identify how the new knowledge can be used to address the more fundamental gaps in their domain and hence produce the most impactful innovations. This creative advantage of specialists means that they are likely to be particularly active (and sought-after) collaborators in fast-evolving knowledge domains.

Testing this argument empirically is challenging because the evolution of knowledge domains is endogenous to the knowledge creation process. In other words, it is difficult to separate

[2] We acknowledge that different knowledge domains evolve at different paces. Our conjecture in this paper is that as the pace of knowledge advancement increases, the relative creative advantage of specialists over generalists gradually grows (or the relative disadvantage at which they find themselves in the slow-moving environment gradually dissipates and they potentially gain an advantage). In this article, we do not intend to investigate the pace of knowledge advancement at which specialists show absolute superior creative advantage over generalists. Rather, we aim to complement the past literature by pointing out that in certain knowledge domains, where the frontier of knowledge moves faster than a certain threshold, specialists would show a creative advantage over their generalist counterparts. We call these knowledge domains fast-evolving. Hence, the dichotomy between fast-evolving knowledge domains and relatively stable ones in our paper is a simplified distinction we make for our theoretical reasoning and our empirical exercise. One should, however, note that what constitutes a fast-pacing knowledge domain in terms of the pace of knowledge advancement may change from one domain to another.

the evolution of a knowledge domain from the efforts and actions of specialist and generalist innovators in that domain. We address this challenge by taking advantage of a natural experiment. Namely, we exploit an unexpected movement of the knowledge frontier in theoretical mathematics that is exogenous to the knowledge creation behavior of innovators leading to the event. Following the unexpected collapse of the Soviet Union in 1989, the ground-breaking work of Soviet theoretical mathematicians suddenly became available to non-Soviet scientists. We take advantage of the fact that Soviet mathematicians were ahead of their non-Soviet colleagues in some domains (e.g., partial differential equations and operator theory) but less so in others (e.g., abstract harmonic analysis and sequences, series, summability) and compare the output of specialists and generalists between domains of theoretical mathematics most and least influenced by the Soviet shock (Agrawal, Goldfarb, and Teodoridis, 2016). The event provides a rare opportunity to analyze the causal effect of specialization for creative performance in evolving knowledge domains.

SPECIALIZATION, COGNITION AND COMPETITION

Specialization as a Hurdle to Creativity

There is no shortage of studies arguing that specialization hinders creativity and innovation. Historians of science and culture have described many cases in which outsiders appear to have benefited from their lack of specialization to produce breakthroughs in such wide-ranging fields as physics (Kuhn 1970) neuroendocrinology (Latour and Woolgar, 1979), astronomy (Edge, 1977), psychoanalysis (McLaughlin, 2001), jazz (Kirschbaum and Vasconcelos, 2006), and painting (Sgourev, 2013). The underlying idea is a familiar one. Lack of specialization can foster creativity because it allows individuals to borrow unorthodox ideas and methods from other domains. Some degree of ignorance seems helpful, a phenomenon that Merton and Zuckerman (1972, 519) referred to as “focused naiveté” or “focused ignorance.” For example, it was the telegraphy amateur Alexander Bell, a speech therapy professional, and not his specialist competitor Elisha Gray, who pursued the development of the “talking telegraph.” Bell was not influenced by the domain-specialist understanding of the telegraph as a toy-like device (Hounshell, 1975). Similarly, John Harrison solved the longitude problem by using clockwork rather than by searching the more specialist solution space of astronomy for which the Longitude Board Principal Scientific Advisor Sir Isaac Newton himself advocated (Andrewes, 1996).

Psychologists have also found considerable support for the idea that specialization can hinder creativity. Specifically, they have 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; Bilalic et al., 2008a: 653). The effect was famously documented in Abraham Luchins' (1942) water-jug experiment. In this short study, participants faced a set of six 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 the former method, but for which a much simpler method also existed. Of the participants who had been exposed to the full set of Einstellung Problems, none found the simpler solution. In contrast, over 60 percent of participants in the control group identified the simpler solution. In other words, prior experience had led to routinized problem-solving which had 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, e.g., references in Bilalić, McLeod, and Gobet, 2008; Dane, 2010). Specialization in a domain can also lead to the formation of habitual behaviors that are rooted in one's knowledge structure (Aarts, Verplanken, and van Knippenberg, 1998; Aarts and Dijksterhuis, 2000; Murray and Häubl, 2007). While habits can increase efficiency in dealing with routine tasks required in a domain, they can nonetheless slow down the pace of adaptation to new tasks and heuristics. Therefore, specialists may find it more difficult to adapt new methods from outside their domain.

Drawing on both traditions, innovation and creativity scholars have argued that specialists ought to produce relatively incremental contributions, while some of the most novel ideas would stem from outsiders. Cattani and Ferriani, for example, examine the Hollywood film industry and find that "the periphery allows one to explore ideas and information not yet widely shared throughout the network" (Cattani and Ferriani, 2008: 827). In a very different setting, Jeppesen and Lakhani (2010) find that scientists that were distant from the focal field of a problem proposed significantly better solutions than insiders did. Others have documented similar insights in the disk drive industry (Audia and Goncalo, 2007) and the computer industry (Bayus, 2012).

Together, these theoretical arguments and empirical findings suggest that specialists tend to exclusively draw upon the knowledge components and heuristics in their domain of specialty. The use and reuse of the same set of familiar knowledge components can, in turn, hinder the process of creative idea generation and innovation. Innovation is most often a process of knowledge recombination (Fleming, 2001). New ideas are essentially combinations of previously disconnected knowledge components. Thus, the heavy reliance of specialists on a narrow domain of knowledge can cause a creative barrier once most of the useful recombinations based on established knowledge in the domain are exhaustively examined (Fleming, 2001). Gradually, new ideas drawn from the same set of knowledge components in a domain become simple extensions of previously established ideas, leading to, at best, incremental innovations (Audia and Goncalo, 2007).[3]

To avoid the downsides of specialization, creative workers are often encouraged to gain exposure to diverse knowledge (Burt, 2004; Audia and Goncalo, 2007; Fleming, Mingo, and Chen, 2007; Lazear, 2005). Having a diverse knowledge base and access to varied sets of heuristics increases the number of recombination opportunities (Taylor and Greve, 2006; Audia and Goncalo, 2007). A vast literature has accordingly linked creative performance to the ability to bring together various knowledge domains not only at the individual level (Burt 2004; Lingo and O'Mahony 2010) but also at the level of teams (Hargadon and Sutton, 1997; Fleming, Mingo, and Chen, 2007), scientific fields (Azoulay, Fons-Rosen, and Zivin, 2015), organizations (March, 1991; Rosenkopf and Nerkar, 2001), cities (Johansson 2004), and regions (Saxenian, 1994). The mechanism is straightforward. Individuals exposed to more diverse knowledge components and heuristics will be less likely to get "trapped" in specific thought patterns. Moreover, the opportunity to bring knowledge from one domain to another opens the door to entirely new creative recombinations. As a result, individuals with knowledge of multiple domains are likely to be more

[3] Research also suggests that there are significant barriers in resolving the entrenchment problem through assembling a diverse team of specialists and facilitating "coordinated exploration" among them (Knudsen and Srikanth, 2014). Using a simulation model, Knudsen and Srikanth (2014) highlight how the incompatibility of mental models between specialists from different backgrounds can lead to conflicts, mutual confusion, joint myopia, and eventually convergence around objectives that are "minimally acceptable for all team members." Several empirical studies have documented the negative impact of intense communication between specialists on innovation outcomes within new product development teams (Tyre and Hauptman, 1992; Hauptman and Hirji, 1996; Song and Montoya-Weiss, 1998; Song, Thieme, and Xie, 1998).

creative than those with expertise in a narrow domain of specialization (e.g., Burt, 2004; Taylor and Greve, 2006).

These arguments implicitly assume that the set of knowledge components in a domain is relatively stable. Otherwise, the documented perils of specialization would not hold, as new knowledge components will always be available to specialists. As it stands, access to components from multiple domains is necessary to identify new recombinations. However, the stability of knowledge components varies significantly across domains and over time. Knowledge domains may undergo periods of incremental or rapid expansion due to scientific and technological discoveries. For example, scientific and technological change constantly increase the set of potential recombinations resulting from regular research as well as from larger-scale initiatives such as the Human Genome Project (Collins and McKusick, 2001) and open-source communities (e.g., Lakhani and von Hippel, 2003). Companies might also try to accelerate innovation through the creation of patent commons (Wen, Ceccagnoli, and Forman, 2015; Vakili, 2016). Similarly, policy-makers attempt to affect the evolution of knowledge domains by increasing access to knowledge components, an interest exemplified by the current debates around the “Fair Access to Science and Technology Research (FASTR) Act,” which requires the open disclosure of peer-reviewed research funded by U.S. science agencies (Kaiser, 2015).

Specialization in Evolving Domains

We argue that specialists have creative advantage over generalists when knowledge domains evolve quickly and substantial new knowledge components become available. The same set of knowledge assets and taken-for-granted thought processes that might blind specialists to exploring cross-domain recombinations could, in fact, enable them to more effectively absorb the newly uncovered knowledge in their domains of specialty, to identify new opportunities for scientific and technological development, and to act upon these opportunities efficiently. These advantages stem from specialists’ deeper understanding of their areas as well as from their familiarity with domain-specific problem-solving heuristics.

Specialists have a more thorough and nuanced grasp of the available knowledge in their domain of specialty. For one, specialists can recall larger amounts of domain-specific knowledge more effectively. For example, Chase and Simon (1973) show that chess masters can recall the exact position of every piece on a chessboard by observing the board very briefly. Specialists also have a more sophisticated appreciation of the different attributes of each component in their knowledge domain as well as the relationships between those components (Dane, 2010). Overall, specialists have a more complete and more accurate perspective of the knowledge landscape in their field. They are better equipped to see and understand the various knowledge gaps in their domain and to recognize the opportunities to address them. Thus, when exposed to new knowledge components in their domain of expertise, specialists are more likely to see how those might be productively recombined.

Specialists are also likely to be particularly well equipped with a variety of domain-relevant problem-solving heuristics. Kuhn (1970) notes that problem-solving is a fundamental aspect of scientific training because most scientific problems are solved by grouping objects and situations into similar sets. Through exposure to a large number of examples, specialists increase their ability to efficiently solve puzzles that are specific to their knowledge domain. In experiments, psychologists have found that physics and mathematics specialists use distinctive problem-solving techniques (Larkin et al., 1980; Sweller, Mawer, and Ward, 1983). In line with Kuhn's argument, Chi, Glaser and Rees (1982) show that expert physicists tend to categorize domain-specific physics problems based on physics principles, whereas non-experts categorize based on 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. Overall, then, specialists are likely to have access to a larger set of domain-specific problem-solving heuristics and might therefore be, on average, better at picking the more appropriate and effective strategy to tackle a given problem in their domain specialty.

Since they are endowed with a deeper understanding of domain-specific knowledge and heuristics, specialists are likely to be particularly well equipped to take advantage of changes in their knowledge domain. This favorable position ought to manifest itself in more than one way. Most directly, specialists are likely to identify the implications of the newly uncovered knowledge more quickly and, therefore, to take advantage of them first. This means that the emergence of domain-specific knowledge is likely to disproportionally increase the creative performance of specialists compared to their generalist colleagues. We therefore hypothesize:

H1: In fast-evolving knowledge domains, specialists will produce more creative output than their generalist colleagues.

Studies have pointed to the importance of distinguishing average creative performance from the propensity to generate breakthroughs (e.g., Fleming, 2001; Audia and Goncalo, 2007). The idea underlying this distinction is that average performance can be driven by a large number of incremental improvements but that breakthroughs can result from a distinct and more divergent process. Again, one underlying assumption of this line of thought is that knowledge domains are relatively stable. In such cases, breakthroughs will emerge only as a result of taking a divergent and original approach.

Here we argue that in fast-evolving knowledge domains, breakthroughs are not necessarily divergent or even original. Considerable evidence exists that many breakthroughs are anticipated or "in the air," a phenomenon leading to scientific and technological races as well as simultaneous discoveries and inventions (Merton, 1957; Lamb and Easton, 1984; Gladwell, 2008). In fast-evolving knowledge domains, specialists have therefore an advantage at producing breakthroughs in that they are likely to quickly identify the new recombination opportunities due to the emergence of new knowledge in their domain, and better ascertain the most impactful ones. A deep understanding of the fundamental gaps and limitations in specialists' knowledge domains helps them tackle the more central questions in their fields and hence to produce more creative breakthroughs. As such, we also hypothesize:

H2: In fast-evolving knowledge domains, specialists will produce more breakthroughs than their generalist colleagues.

Finally, the fact that specialists enjoy relatively high levels of creative opportunity in fast-evolving domains is likely to influence their collaborative behavior as well. Collaboration for creative tasks in general, and in sciences in particular, is not random. Individuals strategically choose their collaborators (Bikard, Murray, and Gans, 2015). Creative opportunities are likely to shape not only scientists' creative performance but also their decision to collaborate. In fast-changing knowledge domains, specialists are likely to better understand the evolving frontier and to therefore be better able to bring together groups of researchers to pursue such opportunities. In addition, others are likely to respond to the growing advantage of specialists by attempting to work with them. For example, IBM's breakthrough development of copper interconnects to replace aluminum ones created a large demand for collaboration with IBM scientists (Lim, 2009). In other words, in their attempts to adapt to the new complexity of the changing knowledge landscape, knowledge workers are likely to disproportionately seek out specialists.

It follows that the rapid movement of the knowledge frontier in a domain is likely associated with disproportionate collaborative opportunities for specialists, a phenomenon consistent with the "burden of knowledge" hypothesis (Jones, 2009; Agrawal, Goldfarb, and Teodoridis, 2016).[4] Although these differences in opportunities cannot be observed directly, they can be measured indirectly by observing the collaborative behavior of creative workers. We therefore expect:

H3: In fast-evolving knowledge domains, specialists will increase collaboration more than their generalist colleagues.

METHODS

Empirical Setting

To test these predictions, we focused on the field of theoretical mathematics and the publication output of mathematics scientists. Similar to prior literature using patents as measures of creative output produced by knowledge workers in R&D labs (Fleming, Mingo, and Chen, 2007; Audia and Goncalo, 2007; Katila, Rosenberger, and Eisenhardt, 2008; Kaplan and Vakili, 2015), we used scientific publications as measures of creative output produced by scientists.[5] Moreover, we exploited a natural experiment – the collapse of the Soviet Union 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 an evolving knowledge domain.

[4] The "burden of knowledge" hypothesis asserts that successive generations of innovators face an increasing education burden due to the advancing knowledge frontier. This, in turn, requires innovators to specialize in increasingly narrower niches and to work more collaboratively over time.
[5] Scholars in other fields have previously used publication output as measures of innovation and creative output in academia. For example, see Jones and Weinberg (2011), Uzzi et al. (2013), and Furman and Stern (2011).

First, the unexpected, exogenous release of new knowledge in certain areas of theoretical mathematics due to the collapse of the Soviet Union enabled us to control for the endogenous link between the activity of creative workers and the evolution of a knowledge domain. Specifically, our empirical strategy relied on the assertion that the collapse of the Soviet Union caused an outward shift in the knowledge frontier in theoretical mathematics and that it did so more for some subfields than for others. We based this claim on three main observations: 1) the Soviet Union's effect on the knowledge frontier in theoretical mathematics was significant for scientific advancements in mathematics; 2) the Soviet Union's effect on the knowledge frontier was greater in some subfields than in others, and the reason for this differential impact is not correlated with active efforts to focus advancements on areas of research away from the rest of the world; and 3) Soviet mathematicians kept their advancements secret from the outside world, for reasons related to the Soviet government's rulings. [6] 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 it offered financial rewards superior to those of many other occupations. However, although Soviet mathematics was strong across the entire spectrum of mathematics, Soviet mathematicians made the greatest advancements in some subfields more than in others (Graham, 1993). Moreover, these differences reflect historical path dependency. Specifically, some subfields of theoretical mathematics built on strong mentorship from the early 1900s and thus continued to attract bright minds later on (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. Furthermore, Soviet knowledge in theoretical mathematics was kept secret from the outside world due to the Communist government's rules and regulations. The Soviet government kept strict control on international travel. Academics who wished to attend foreign conferences had to undergo a stringent and lengthy approval process, and many researchers were blacklisted because of "tainted" backgrounds. The few approvals granted were typically for travel in Eastern Europe (Ganguli, 2014). Additionally, Soviet researchers were prevented from publishing their findings, traveling

[6] See Agrawal, Goldfarb, and Teodoridis (2016) for a more detailed discussion of the impact of the Iron Curtain on theoretical mathematics. In addition to concerns of exogeneity of this event that the authors address, their paper also provides extensive evidence of the event as a shock to the knowledge frontier and not one driven by competition in labor markets.

to conferences, communicating or collaborating with non-Soviets, or even accessing non-Soviet references. As such, Soviet advancements in mathematics remained relatively unknown to the outside world until the collapse of the Soviet Union (Graham and Dezhina, 2008), when they were suddenly made available.[7]

Second, 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 time. 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.

Third, we observed the specialization level of mathematicians in our sample based on a manual detailed categorization of research output provided by the MR Database. Specifically, we relied on the careful and exhaustive work of the Mathematical Reviews division, which classifies each paper in mathematics using Mathematics Subject Classification (MSC) codes. The MSC schema is internationally recognized and facilitates targeted searches on research subjects across all subfields of mathematics. The MR team assigns precisely one primary MSC code to each academic publication uploaded to the MR Database. There are 33 codes covering theoretical mathematics (Table 1). Using the MSC codes assigned to each paper, we can measure the specialization degree of each individual mathematician at a given time.

– Insert table 1 about here –

Fourth, compared to other knowledge domains, the field of theoretical mathematics is a conservative setting in which to test our predictions. The explicit, codified nature of knowledge in theoretical mathematics lowers the barriers to knowledge acquisition and use for generalists. In domains with larger shares of tacit knowledge, such as medical sciences, it is arguably more

[7] 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.*

difficult for generalists and non-experts to acquire and use the tacit knowledge embedded within specialists' circles.

Last, the field of theoretical mathematics plays a fundamental role in knowledge and technological progress across a wide range of domains. The examples are numerous. Wavelet and Fourier transforms are widely used in electronics, computer graphics, and medical equipment such as MRI machines. Algebraic Toplogy is used extensively in data mining and processing. Number theory, and particularly the theory of prime numbers, has immensely influenced the 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. To put it simply, theoretical mathematics provides the abstract foundation and structure to formulate and understand the physical world around us. 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

As noted, we collected our data from the Mathematical Reviews Database, the most comprehensive database on academic publications in the field of mathematics. The 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, between 1980 and 2000.

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 includes the year of publication, the MSC classification code, the full set of authors per academic publication, and the number of academic citations received from subsequent publications. Next, we re-arranged 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, since these individuals were already at the frontier of knowledge, and focused on all other mathematicians, since their group experienced the frontier advancement. We also excluded all publications with at least one Soviet author to make sure the 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 collapse of the Soviet Union, namely between 1980 and 1989.[8] The choice of a minimum of four publications helped us to carefully separate specialists from generalists in our sample, and to ensure that our results are not driven by unproductive individuals classified as specialists due to their low number of publications. For example, individuals with one publication would otherwise be automatically classified as specialists. However, 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 in the following sections, we identified all individuals that could be cleanly categorized as either specialist or 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 collapse of the Soviet Union. As we discuss below, there are some significant differences between specialists and generalists in terms of productivity in the years prior to the collapse of the Soviet Union. This is not surprising, since 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. As such, to ensure that our results are not biased due to 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 collapse of the Soviet Union. 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 record in the pre-collapse period.[9] The matched sample contains data on 4,076 mathematicians, of which

[8] 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.
[9] 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, as predicted in hypothesis 1, we used their quality-adjusted count of publications per year. More specifically, 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 from subsequent publications to construct a citation-weighted count of publications per mathematician per year. 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, her quality-adjusted research output for 1985 is 32. As a robustness check, we also report estimations based on a simple count of publications per year.

Hypothesis 2 predicts the relative advantage of specialists over generalists in producing breakthroughs in areas that experienced a substantial influx of new knowledge due to the collapse of the Soviet Union. To test this hypothesis, we constructed a measure of breakthrough output based on the citation distribution of all publications in our sample. Following past research (Ahuja and Lampert, 2001; Phene, Fladmoe-Lindquist, and Marsh, 2006; Kaplan and Vakili, 2015), we first coded the publications that belong to the top 5 percent of highly cited publications in any given year as breakthroughs. Next, we count the number of breakthroughs for each individual mathematician in any given year to construct an individual-level measure of breakthrough output per year. As 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.

Finally, to test hypothesis 3, we constructed two measures to capture the number of collaborators for each mathematician per year. The first measure simply counts the total number of co-authors on all publications of an individual in any given year. The second measure captures the number of unique co-authors with whom each mathematician worked in a given year.

Independent Variables

We used three indicators (and their interactions) as main independent variables in all estimations. The first variable, $Specialized_{i,t}$, determines whether a mathematician in our sample is a specialist or generalist at the time of the collapse of the Soviet Union. 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 of the Soviet Union (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 (Table 1) 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:

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

By construction, the higher the value of $DiversificationIndex_i$, the greater the breadth of areas in which mathematician i published before the collapse of the Soviet Union. The diversification measure is greater than or equal to 0 and never reaches 1. The highest possible value of the diversification index is 0.83 and would characterize 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 exclusively published in one subfield of theoretical mathematics. 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 0.29, while an equally productive colleague who published all her papers in one subfield of theoretical mathematics would have a diversification index of 0. In our sample, the highest diversification index is 0.531 and the lowest is 0. In our main specification, we define generalists as mathematicians with a diversification index in the top 10 percent of the distribution (above 0.290) and specialists as those with a diversification index of 0 (i.e., those who only published in one subfield). Our results remain robust to considering alternative cut-off points that continue to capture the variation in the breadth of expertise.

The second variable, $SovietImpact_i$, determines the degree to which each mathematician in our sample was affected by the collapse of the Soviet Union. The variable separates mathematicians who experienced a substantial movement of the knowledge frontier in their areas from others. 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 collapse of the Soviet Union. Table 1 lists the 33 subfields and their rank. Based on this ranking, 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 collapse of the Soviet Union. The higher the percentage of one's publications in subfields where Soviets made higher 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 takes into account the full publication portfolio during the period before the collapse (1980–1989). For example, a mathematician who published all his papers in "Integral Equations," the most affected subfield of theoretical mathematics, would have a Soviet impact index of 1. If he were to publish all his papers in "Fourier Analysis," the second most affected subfield of theoretical mathematics, his Soviet impact index would be 0.5. And if he were to publish half his work in "Integral Equations" and half in "Fourier Analysis," his Soviet index impact would be 0.75. In our sample, the minimum value of the Soviet impact index is 0.030, the maximum value is 1, the mean is 0.108, and the standard deviation is 0.112. 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 different cut-off values to capture the variation in Soviet impact.[10]

[10] We chose to 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

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 their 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 their cumulative number of publications (since 1980). The variable is logged to account for its skewed distribution. Last, we controlled for the non-linear effect of age by including an age-squared term in all estimations.[11] Since 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 and collaboration rates of specialists and generalists affected by the forward movement of the knowledge frontier in theoretical mathematics due to the collapse of the Soviet Union. 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. Specifically, we compared the change in creative output and collaboration rates of specialists with that of generalists in areas most affected by the fall of the Soviet Union, using the difference in areas less affected as the baseline. Formally, we estimated:

$$\begin{aligned} DV_{i,t} = f\big(&\beta_1.Specialist_i.SovietImpact_i.AfterSovietCollapse_t \\ &+ \beta_2.SovietImpact_i.AfterSovietCollapse_t + \beta_3.Specialist_i.AfterSovietCollapse_t \\ &+ C_{i,t} + I_i + \gamma_t + \varepsilon_{i,t}\big) \end{aligned}$$

magnitude of the estimation results. Our findings are robust to using a continuous measure and to analyzing the marginal effects at different parts of the distribution.

[11] Note that the inclusion of both individual fixed effects and year fixed effects automatically controls for the linear effect of individual age.

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 since their values are fixed at the individual level. Similarly, $AfterSovietCollapse_t$ is not included independently since its effect is absorbed by the year fixed effects.

β_1 is the main coefficient of interest and captures the differential performance of specialists relative to generalists in areas of theoretical mathematics most affected by the collapse of the Soviet Union, using the difference in areas less affected as the baseline (i.e., counterfactual). β_2 captures the change in the outcome trend of generalists who were active in areas affected by the Soviet collapse compared to generalists whose research was predominantly outside the affected areas. β_3 captures the difference between post-Soviet outcome trends of specialists and generalists whose research was primarily outside the areas affected by the Soviet collapse. While we were primarily interested in β_1 to test our three predictions, β_3 can provide some insight into the differential performance of specialists relative to generalists in areas that did not experience a meaningful movement in knowledge frontier.

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 in all estimations. The estimator is consistent in the presence of heteroskedasticity and over-dispersion of the dependent variable (Silva and Tenreyro 2006).

RESULTS

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 0.7 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, some 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 in the full sample due to 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 2 about here –

Table 3 details the differences between specialists and generalists on the key dimensions of interest for the period before the collapse of the Soviet Union. Panel A shows the differences in the full sample, and Panel B reports them in the matched sample.

– Insert table 3 about here –

In the full sample (Panel A) there are almost twice as many specialists as generalists. This is not surprising given the graduate training of these individuals 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 0.2 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 take-away is that the differences between specialists and generalists in the full sample disappear once we restrict our sample to the CEM one. The only somewhat persistent difference is in the number of unique collaborators, where generalists still have a more diverse set of collaborators. We address this aspect by estimating changes in collaboration using the two variables described previously, rather than relying only on a measure of unique collaborators. Last, due to our strict one-to-one matching, the number 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.

Table 4 shows estimation results for the change in the creative output of specialists versus generalists in areas most affected by the collapse of the Soviet Union. Model 1 shows results for the citation-weighted number of publications, and model 2 shows results for the simple count of publications. The estimated β_2 in model 1 suggests that, in mathematics areas that did not experience a substantial movement in knowledge frontier, specialists produced approximately 25 percent fewer citation-weighted publications per year than generalists did in years after 1989. These results are generally consistent with the assertion that generalists have a creative advantage over specialists when a knowledge domain remains stable over time. In contrast, using the change in the differential performance of specialists versus generalists in less affected areas as the baseline, the estimated β_1 suggests that in the most affected area of mathematics, specialists increased their citation-weighted publication output relative to generalists by approximately 83 percent in years after 1989. It is equivalent to receiving approximately 4 more citations per year during the post-Soviet period. This creative advantage of specialists is partly attributed to the decline in the performance of generalists in the most affected areas after the Soviet collapse, as suggested by the negative and significant β_3. However, specialists also increased their performance substantially. Compared to specialists in less affected areas, specialists in the most affected areas increased their performance by approximately 19 percent after the Soviet collapse.

– Insert table 4 about here --

Model 2 of table 4 shows estimation results using count of publications as the dependent variable. The interpretation of results is similar to those reported for model 1. However, the coefficients are smaller, which suggests that the relative increase in creative performance of specialists in affected areas after the Soviet collapse is driven partly by the increase in the quantity of their creative output and partly by the increase in the average quality of their creative output (measured by the number of citations to their publications). The estimated β_2 suggests a significant 7 percent relative decline in the number of publications by specialists in unaffected areas. In contrast, there is a relative increase of approximately 37 percent in the number of publications by specialists over generalists in math areas most affected by the Soviet collapse, using the change in the differential performance of specialists relative to generalists in less affected areas as a baseline.

The 37 percent increase is equivalent to approximately one extra publication every three year. Again, part of this publication premium is driven by a decline in generalists' number of publications in the most affected areas, and part of it is driven by an increase in specialists' publication output. Compared to specialists in the less affected areas of mathematics, specialists in the most affected areas increased their publication count by approximately 6 percent after the fall of the Iron Curtain.

Models 3 and 4 report the analog estimation results for the matched sample. The estimated coefficients are in line with those for the full sample and depict trends similar to those described above. The coefficients are slightly larger, suggesting a larger creative premium for specialists over generalists in areas that experienced a movement in the knowledge frontier after the Soviet collapse. Overall, the results in table 4 are consistent with our prediction in hypothesis 1 and confirm the creative advantage of specialists in areas with an evolving knowledge frontier.

One potential concern with these interpretations is that the observed change in the differential performance of specialists relative to generalists in the most and least affected areas 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 compared to generalists in years before and after the Soviet collapse in 1989. In particular, we used the pre-1984 performance difference of specialists and generalists in the non-affected areas as the baseline and examined the change in citation-weighted output of specialists versus generalists in the most affected areas over six periods: 1984–1986, 1987–1989, 1990–1992, 1993–1995, 1996–1998, and 1999–2000. 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. Using three-year periods thus helps minimize the noise due to mathematicians' infrequent research output, a norm of the field of theoretical mathematics. In testing for pre-trends, we used the matched sample to ensure maximum comparability between specialists and generalists in the period before the collapse. If changes in the differential performance of specialists and generalists indeed pre-date the Soviet collapse, we should observe these trends in the 1984–1986 and 1996–1998 periods. The results are depicted in figure 1. The estimates suggest that the change in the differential performance of specialists and generalists in the most affected areas began about two years after the fall of the Iron Curtain. The delay is in line with the publication pace in theoretical mathematics and includes the time it takes for specialists to acquire and absorb new knowledge, identify new opportunities based on that knowledge, exploit the identified opportunities, and undergo the publication process. There are no indications of pre-Soviet collapse trends, confirming that the estimated changes in table 1 are attributed to years following the fall of the Soviet Union.

– Insert figure 1 about here –

To test our second hypothesis, that specialists also produce more breakthroughs in areas with a moving knowledge frontier, in table 5 we display results for the differential rate of breakthrough production between specialists and generalists. The estimates are in line with those reported in table 1 for overall creative output. The estimated β_2 in model 1 indicates a significant decline of about 25 percent in the number of breakthroughs – defined as publications in the top 5 percent of cited output – produced by specialists relative to that produced by generalists in areas less affected by the Soviet collapse. Again, the results are consistent with the idea that generalists have more opportunities to generate breakthroughs in stable knowledge domains due to their more diverse knowledge sources. In contrast, using the change in the differential performance of specialists and generalists in less affected areas as the baseline, the estimated β_1 suggests that specialists produced on average 74 percent more breakthroughs than generalists in areas most affected by the Soviet collapse, post-1989. In other words, specialists were almost twice more likely than generalists to generate a breakthrough after the Soviet collapse in areas where Soviet mathematicians had the greatest impact on the knowledge frontier. Again, the negative, though not statistically significant, β_3 suggests that part of this differential breakthrough output by specialists can be explained by a decline in the breakthrough output of generalists in the most affected areas. Nevertheless, the difference is largely attributed to an increase in the performance of specialists. The results hold when replacing our dependant variable with the alternative measure of breakthroughs based on the top 10 percent cited publications. The results also persist in the matched sample, with the estimates suggesting a slightly larger relative increase in the breakthrough output of specialists compared to generalists in mathematics areas most affected by the Soviet collapse.

– Insert table 5 about here –

As before, to examine the timing of effects, we estimated the relative change in breakthrough output of specialists versus generalists in the most affected areas over the six periods defined, using the change in the relative performance of specialists and generalists in least affected areas as the baseline. The steps we took are the same as those taken to construct figure 1. We present results in figure 2 and figure 3 for breakthroughs defined as top 5 percent cited and top 10 percent cited publications, respectively. As before, we found no evidence of pre-trends, but only of an increase after the collapse of the Soviet Union. Compared to figure 1, the estimated coefficients for the periods after the Soviet collapse have larger standard errors. To this point, we draw attention to the rare nature of breakthrough output and hence the smaller variance in our dependent variables for this particular set of estimations.

– Insert figures 2 and 3 about here –

To test our third hypothesis, in table 6 we present results for the change in the collaboration rates of specialists versus generalists after the fall of the Soviet Union. The estimated β_2 in models 1 and 2 suggests a relative decline of 7 percent and 10 percent in specialists' number of collaborators and specialists' number of unique collaborators, respectively, compared to generalists in the less affected areas. In the most affected areas, however, specialists' total number of collaborators and unique collaborators increased by 40 percent and 53 percent respectively compared to generalists, using their differential change in the non-affected areas as the baseline. The change is equivalent to approximately one additional new collaborator every four years. The results are driven partly by generalists working with fewer collaborators in the most affected areas after the Soviet collapse, and partly by specialists working with more and new collaborators in the same areas. Figures 4 and 5 show the timing of change in specialists' collaboration rates relative to generalists in the most affected areas. Again, there is no sign of trends in years before the Soviet collapse.

– Insert table 6 and figures 4 and 5 about here --

In our robustness checks, we tested the sensitivity of our results to varying thresholds to defining specialists and generalists as well as varying thresholds to defining most and least affected areas of theoretical mathematics. Our results remained robust to these alternative threshold choices.

DISCUSSION AND CONCLUSION

While some domain-specific knowledge is undoubtedly important, prior literature has found extensive evidence that specialization can hinder creative work by locking individuals into limited thought patterns, recombination options, and repeat heuristics. Specialists are likely to be embedded in the tradition of their knowledge domains and thus unable to consider unorthodox approaches and to establish divergent and deeply original recombinations.

This paper does not contradict this line of thinking. Rather, we add to it by drawing attention to the fact that its arguments rest on an implicit assumption of relative stability of knowledge domains. In cases where knowledge domains evolve quickly, specialists might be better prepared than their more diversified colleagues – generalists – to identify and exploit opportunities for recombination stemming from advancements made at the knowledge frontier. Knowledge domains advance for various reasons, such as technological breakthroughs or spontaneous surfacing of previously inaccessible knowledge. Specialists are in a better position to make use of the new knowledge in their domains since they have a more thorough understanding of the knowledge landscape and a better grasp of the fundamental debates and gaps in the domain.

We hypothesized and found empirical support for a higher creative performance of specialists in knowledge domains with a fast-evolving frontier. After the collapse of the Soviet Union and the unexpected, and substantial, advancement of several areas of theoretical mathematics, not only did non-Soviet specialists in those areas produce more academic publications, they also produced more breakthroughs than their generalist colleagues did. Furthermore, we found evidence of a disproportionate increase in collaboration rates of specialists relative to their generalist colleagues in areas that experienced substantial advancements of the knowledge frontier. We interpret this result as providing additional support to our main conjecture that specialists are better able to seize creative opportunities in the context of an advancing knowledge frontier. Because of their more advantageous position, they are better able to find and attract collaborators, which means they collaborate more than generalists do in fast-changing knowledge domains as opposed to more stable ones.

In our empirical analysis, we were careful to acknowledge and mitigate, albeit not without limitations, the main identification concerns associated with our theoretical reasoning: the fact that the advancement of knowledge domains is endogenous to the knowledge creation process, in particular with respect to the knowledge creation behavior of generalists and specialists. While the natural experiment of the unexpected collapse of the Soviet Union and the sudden release of knowledge in theoretical mathematics that shifted the knowledge frontier forward in some domains but less so in others provides a rewarding testing bed 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, this is nevertheless a study of one setting. Thus, we make no strong claims of generalizability and invite future research to investigate other knowledge creation areas. We draw attention to the fact that theoretical mathematics is a somewhat conservative field in terms of yearly productivity, collaboration levels, and overall number of academic publications (Agrawal, Goldfarb, and Teodoridis, 2016). Furthermore, theoretical mathematics does not rely on access to research equipment that might influence knowledge output and collaboration decisions (Stephan, 2012; Murray et al., 2016; Teodoridis, 2017).

Second, our study is limited by an empirical ability to acknowledge that levels of specialization and diversity at the individual level might vary over time. Indeed, individuals might change their level of diversification throughout their career. For example, junior researchers might exhibit a higher level of specialization whereas senior individuals might exhibit a higher level of diversification. In our empirical analysis, we controlled for a quadratic effect of age to account for this possibility. However, this approach does not account for the fact that junior specialists might become increasingly diverse as they advance in their career. 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 attempts several theoretical contributions. First, we contribute to the literature on knowledge creation and recombination by providing evidence that challenges the traditional view on specialization in creative work. Prior to this study, a vast body of work focused on highlighting the limitations of specialization in the process of knowledge recombination and advocated for diversity as a more productive avenue that facilitates spanning across various knowledge bases (Hargadon and Sutton, 1997; Uzzi and Spiro, 2005; Audia and Goncalo, 2007;

Cattani and Ferriani, 2008; Jeppesen and Lakhani, 2010). We extend this literature by examining the distinctive advantages of specialization. Specifically, our study suggests that specialists play a crucial role in the knowledge recombination process when the knowledge frontier advances in their domains of expertise. Prior literature suggests that diversification plays an important role in the knowledge creation process during periods of relative stability in knowledge domains.

Second, and related, we propose that the pace of change in creative domains has crucial implications for creative workers and for the organization of knowledge creation. Our study investigated what appears to be a common underlying assumption in many studies of creativity: the relative stability of knowledge domains. Not only do knowledge domains evolve, some evolve faster than others. We propose that the pace of evolution of knowledge domains has important implications for individual workers, as it will benefit some more than others. While our study provides evidence of one context of a fast-advancing knowledge frontier, we hope future research can shed more light on the interaction between the pace of knowledge advancement and the relative advantage of specialists (or generalists) in a knowledge domain. For example, in fast-evolving domains, incentives might encourage specialization; in more stable domains, incentives might encourage diversification of knowledge and expertise. Stated differently, our findings suggest a more complex relationship between specialization and diversification than previously thought, one that is contingent on the varying pace of advancement of knowledge domains.

Third, our study contributes to the discussion on the emergence of breakthroughs. Creative breakthroughs have traditionally been associated with distant search, whereas local search was believed to lead to more incremental improvements (e.g., Fleming, 2001). At its core, this view assumes that truly novel recombinations can only occur by drawing on elements that had not been considered before – that is, those that stem from outside the knowledge domain in question. In contrast, we show that completely new elements can emerge inside a knowledge domain, especially if the latter evolves rapidly. In these situations, breakthroughs are likely to stem from local rather than distant search. Put differently, the promise of local search is likely to vary with the pace of knowledge-domain advancement. In our setting, local search appears to have been a particularly productive strategy in fast-evolving knowledge domains of theoretical mathematics. This could explain why so many breakthroughs are simultaneous discoveries (Merton, 1957). Moreover, it is especially important given the well-documented skewed nature of innovation (e.g., Trajtenberg, 1990), where a handful of highly impactful innovations fuel organizational and economic growth.

Fourth, our study contributes to the literature on collaboration in knowledge creation by contextualizing the common finding that collaboration is associated with high creative performance. While prior studies have argued that the link between collaboration and creativity should be interpreted as a sign that collaboration fosters creativity (e.g., Reagans, Zuckerman, and McEvily, 2004; Fleming, Mingo, and Chen, 2007; Singh and Fleming, 2010), our study points to another interpretation of this link. Specifically, we suggest that, in some cases, collaboration might actually be the result of creative opportunities. In other words, collaboration and creative performance might co-occur not because one drives the other but because both result from the advantageous position of some individuals (here, specialists in the context of evolving knowledge domains). This alternative mechanism is important in that it points to the hitherto under-recognized importance of the uneven distribution of creative opportunities as an unobserved link between collaboration and creativity. At one extreme, if creative workers collaborate when they have better creative opportunities and not the other way around, then prior studies may have overstated the benefits of collaboration for creative performance.

Additionally, our study has implications for the organization of firms' R&D. At one level, it suggests that the composition of the R&D team, with respect to the balance of specialists and generalists, should depend on the pace of innovation in the industry. However, it also suggests that R&D team composition has important implications for the firm's ability to absorb external knowledge (Cohen and Levinthal 1990). Specifically, although specialists might run the risk of always using the same heuristics, they may be better able to absorb external knowledge in their domains of expertise and, in so doing, be likely to help firms navigate fast-changing environments. This is important not only for increasing organizational performance but, at a more fundamental level, for organizational survival. Prior work has emphasized the importance of R&D not only as a product development channel but also as an avenue for firms to gain new knowledge and keep abreast of the competition (e.g., Cockburn and Henderson, 1998; Owen-Smith and Powell, 2004).

Our paper does not argue that specialization is the key to creative performance. Rather, we highlight that specialization has advantages that prior literature might have overlooked. Certainly, generalists have a superior ability to broker ideas across knowledge domains, and this ability is likely to be very advantageous in the more stable settings (e.g., Hargadon and Sutton, 1997). However, many knowledge domains change rapidly, and specialization has crucial benefits in those situations. In other words, our study suggests that specialists and generalists fulfill different and complementary creative roles. In highlighting the distinctive role played by specialists, we hope to enhance our understanding of the drivers of creative performance and invite future research to consider the nuanced relationship between diversification and specialization in creative work.

REFERENCES

Aarts, H., and A. Dijksterhuis. 2000. "Habits as Knowledge Structures: Automaticity in Goal-Directed Behavior." Journal of Personality and Social Psychology, 78 (1): 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 (15): 1355–74.

Agrawal, A., A. Goldfarb, and F. Teodoridis. 2016. "Understanding the Changing Structure of Scientific Inquiry." American Economic Journal: Applied Economics, 8 (1): 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 (6/7): 521–43.

Amabile, T. M. 1983. "The Social Psychology of Creativity: A Componential Conceptualization," Journal of Personality and social Psychology, 45(2): 357-376.

Andrewes, W. J. H., ed. 1996. The Quest for Longitude. Illustrated edition. Cambridge, Mass.: Collection of Historical Scientific Instruments.

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): 1–15.

Azoulay, P., C. Fons-Rosen, and J. S. Graff Zivin. 2015. "Does Science Advance One Funeral at a Time?" Working Paper 21788. National Bureau of Economic Research. http://www.nber.org/papers/w21788.

Azoulay, P., T. Stuart, and Y. Wang. 2013. "Matthew: Effect or Fable?" Management Science, 60 (1): 92–109.

Bayus, B. L. 2012. "Crowdsourcing New Product Ideas over Time: An Analysis of the Dell IdeaStorm Community." Management Science, 59 (1): 226–44.

Benner, M. J., and M. L. Tushman. 2003. "Exploitation, exploration, and process management: The productivity dilemma revisited." Academy of Management Review, 28(2): 238-256.

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 (7):

Bilalić, M., P. McLeod, and F. Gobet. 2008. "Inflexibility of experts—Reality or Myth? Quantifying the Einstellung Effect in Chess Masters." Cognitive Psychology, 56 (2): 73–102.

Blackwell, M., S. Iacus, G. King, and G. Porro. 2009. "Cem: Coarsened Exact Matching in Stata." Stata Journal, 9 (4): 524–46.

Borjas, G. J., and K. B. Doran. 2012. "The Collapse of the Soviet Union and the Productivity of American Mathematicians." The Quarterly Journal of Economics, 127(3): 1143-1203.

Burt, R. S. 2004. "Structural Holes and Good Ideas." American Journal of Sociology, 110 (2): 349–99.

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 (6): 824–44.

Chase, W. G., and H. A. Simon. 1973. "Perception in Chess." Cognitive Psychology, 4 (1): 55–81.

Chi, M. T. H., R.. Glaser, E. Rees, and R. Sternberg. 1982. Advances in the Psychology of Human Intelligence. New York: Psychology Press.

Cockburn, I. M., and R. M. Henderson. 1998. “Absorptive Capacity, Coauthoring Behavior, and the Organization of Research in Drug Discovery.” The Journal of Industrial Economics, 46 (2): 157–82.
Cohen, W. M., and D. A. Levinthal. 1990. “Absorptive Capacity: A New Perspective on Learning and Innovation.” Administrative Science Quarterly, 35 (1): 128–52.
Collins, F. S., and V. A. McKusick. 2001. “Implications of the Human Genome Project for Medical Science.” JAMA 285 (5): 540–44.
Dane, E. 2010. “Reconsidering the Trade-off Between Expertise and Flexibility: A Cognitive Entrenchment Perspective.” Academy of Management Review, 35 (4): 579–603.
Edge, D. 1977. “The Sociology of Innovation in Modern Astronomy.” Quarterly Journal of the Royal Astronomical Society, 18 (3): 326–39.
Fleming, L. 2001. “Recombinant Uncertainty in Technological Search.” Management Science, 47 (1): 117–32.
Fleming, L., S. Mingo, and D. Chen. 2007. “Collaborative Brokerage, Generative Creativity, and Creative Success.” Administrative Science Quarterly, 52 (3): 443–75.
Forman, C., A. Goldfarb, and S. Greenstein. 2016. “Agglomeration of Invention in the Bay Area: Not Just ICT.” American Economic Review, 106 (5): 146–51.
Franzoni, C., G. Scellato, and P. Stephan. 2011. “Changing Incentives to Publish.” Science, 333 (6043): 702–3.
Frensch, P. A., and R. J. Sternberg. 1989. “Expertise and Intelligent Thinking: When Is It Worse to Know Better.” Advances in the Psychology of Human Intelligence, 5: 157–188.
Furman, J., and S. Stern. 2011. “Climbing atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research.” American Economic Review, 101 (5): 1933–63.
Gambardella, A. 1992. “Competitive Advantages from in-House Scientific Research: The US Pharmaceutical Industry in the 1980s.” Research Policy, 21 (5): 391–407.
Ganguli, I. 2014. “Scientific Brain Drain and Human Capital Formation after the End of the Soviet Union.” International Migration, 52 (5): 95–110.
Gladwell, M. 2008. “In the Air.” The New Yorker, May 12. http://www.newyorker.com/magazine/2008/05/12/in-the-air.
Graham, L. R. 1993. Science in Russia and the Soviet Union: A Short History. Cambridge: University Press.
Graham, L. R., and I. Dezhina. 2008. Science in the New Russia: Crisis, Aid, Reform. Indiana: University Press.
Hargadon, A., and R. Sutton. 1997. “Technology Brokering and Innovation in a Product Development Firm.” Administrative Science Quarterly, 42 (4): 716–49.
Hauptman, O., and K. K. Hirji. "The influence of process concurrency on project outcomes in product development: An empirical study of cross-functional teams." IEEE Transactions on Engineering Management, 43(2): 153-164.
Hounshell, D. A. 1975. “Elisha Gray and the Telephone: On the Disadvantages of Being an Expert.” Technology and Culture, 16 (2): 133–61.
Iacus, S. M., King, G., & Porro, G. 2011a. “Causal inference without balance checking: Coarsened exact matching.” Political Analysis, 20(1): 1-24.
Jeppesen, L. B., and K. R. Lakhani. 2010. “Marginality and Problem-Solving Effectiveness in Broadcast Search.” Organization Science, 21 (5): 1016–33.
Johansson, F. 2004. The Medici Effect: Breakthrough Insights at the Intersection of Ideas, Concepts, and Cultures. Cambridge: Harvard Business Press.

Jones, B. 2009. “The Burden of Knowledge and the ‘Death of the Renaissance Man’: Is Innovation Getting Harder?” Review of Economic Studies, 76 (1): 283–317.
Jones, B. F., and B. A. Weinberg. 2011. "Age dynamics in scientific creativity." Proceedings of the National Academy of Sciences 108(47): 18910-18914.
Kaiser, J. 2015. “Senate Panel Approves Public Access Bill.” Science. July 29. http://www.sciencemag.org/news/2015/07/senate-panel-approves-public-access-bill.
Kaplan, S., and K. Vakili. 2015. “The Double-Edged Sword of Recombination in Breakthrough Innovation.” Strategic Management Journal, 36 (10): 1435–1457.
Katila, R., J. D. Rosenberger, and K. M. Eisenhardt. 2008. “Swimming with Sharks: Technology Ventures, Defense Mechanisms and Corporate Relationships.” Administrative Science Quarterly, 53 (2): 295–332.
Kirschbaum, C., and F. Vasconcelos. 2006. “Jazz: Structural Changes and Identity Creation in Cultural Movements.” Advances in Organization Studies, 20: 228.
Klincewicz, K. 2016. “The Emergent Dynamics of a Technological Research Topic: The Case of Graphene.” Scientometrics, 106 (1): 319–345.
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(3): 409-441.
Kuhn, T. S. 1970. The Structure of Scientific Revolutions. Chicago: University of Chicago Press.
Lakhani, K. R, and E. von Hippel. 2003. “How Open Source Software Works: ‘free’ User-to-User Assistance.” Research Policy, 32 (6): 923–943.
Lamb, D., and S. M. Easton. 1984. Multiple Discovery: The Pattern of Scientific Progress. Avebury Publishing.
Larkin, J., J. McDermott, D. P. Simon, and H. A. Simon. 1980. “Expert and Novice Performance in Solving Physics Problems.” Science, 208 (4450): 1335–42.
Latour, B., and S. Woolgar. 1979. Laboratory Life: The Construction of Scientific Facts. Princeton, NJ: Princeton University Press.
Lazear, E. P. 2005. “Entrepreneurship.” Journal of Labor Economics, 23 (4): 649–80.
Lim, K. 2009. “The Many Faces of Absorptive Capacity: Spillovers of Copper Interconnect Technology for Semiconductor Chips.” Industrial and Corporate Change, 18 (6): 1249-1284.
Lingo, E. L., and S. O’Mahony. 2010. “Nexus Work: Brokerage on Creative Projects.” Administrative Science Quarterly, 55 (1): 47–81.
Luchins, A. S. 1942. “Mechanization in Problem Solving: The Effect of Einstellung.” Psychological Monographs, 54 (6): i-95.
March, James G. 1991. “Exploration and Exploitation in Organizational Learning.” Organization Science, 2 (1): 71–87.
Marx, M., J. S. Gans, and D. H. Hsu. 2014. “Dynamic Commercialization Strategies for Disruptive Technologies: Evidence from the Speech Recognition Industry.” Management Science, 60 (12): 3103–3123.
McLaughlin, N.. 2001. “Optimal Marginality: Innovation and Orthodoxy in Fromm’s Revision of Psychoanalysis.” The Sociological Quarterly, 42 (2): 271–288.
Merton, R. K. 1957. “Priorities in Scientific Discovery: A Chapter in the Sociology of Science.” American Sociological Review, 22 (6): 635–659.
Mokyr, J. 2002. The Gifts of Athena: Historical Origins of the Knowledge Economy. Princeton, NJ: Princeton University Press.

Murray, F., P. Aghion, M. Dewatripont, J. Kolev, and S. Stern. 2016. "Of Mice and Academics: Examining the Effect of Openness on Innovation." American Economic Journal: Economic Policy 8 (1): 212–52.
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 (1): 77–88.
Owen-Smith, J., and W. W. Powell. 2004. "Knowledge Networks as Channels and Conduits: The Effects of Spillovers in the Boston Biotechnology Community." Organization Science, 15 (1): 5–21.
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 (4): 369–388.
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 (1): 101–33.
Rosenkopf, L., and A. Nerkar. 2001. "Beyond Local Search: Boundary-Spanning, Exploration, and Impact in the Optical Disk Industry." Strategic Management Journal, 22 (4): 287–306.
Russo, E. "Special Report: The birth of biotechnology." Nature, 421(6921): 456-457.
Samelson, F. 1977. "World War I Intelligence Testing and the Development of Psychology." Journal of the History of the Behavioral Sciences, 13 (3): 274–82.
Saxenian, A. 1994. Regional Advantage: Culture and Competition in Silicon Valley and Route 128. Harvard University Press.
Sgourev, S. V. 2013. "How Paris Gave Rise to Cubism (and Picasso): Ambiguity and Fragmentation in Radical Innovation." Organization Science, 24 (6): 1601–17.
Silva, J. M. C. Santos, and S. Tenreyro. 2006. "The Log of Gravity." Review of Economics and Statistics, 88 (4): 641–58.
Singh, J., and L. Fleming. 2010. "Lone Inventors as Sources of Breakthroughs: Myth or Reality?" Management Science, 56 (1): 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(2): 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(4): 289-303.
Stephan, P.. 2012. How Economics Shapes Science. Cambridge, Massachusetts: Harvard University Press.
Sweller, J., R. F. Mawer, and M. R. Ward. 1983. "Development of Expertise in Mathematical Problem Solving." Journal of Experimental Psychology: General, 112 (4): 639–661.
Taylor, A., and H. R. Greve. 2006. "Superman or the Fantastic Four? Knowledge Combination and Experience in Innovative Teams." Academy of Management Journal, 49 (4): 723.
Teodoridis, F. 2017. "Understanding Team Knowledge Production: The Interrelated Roles of Technology and Expertise." Working Paper.
Tikhomirov, V. M. 2007. "On Moscow Mathematics-Then and Now." Golden Years of Moscow Mathematics, Second Edition, Co-Publication of the AMS and the London Mathematical Society.

Trajtenberg, M. 1990. “A Penny for Your Quotes: Patent Citations and the Value of Innovations.” The RAND Journal of Economics, 21 (1): 172–187.
Tyre, M. J., and O. Hauptman. 1992. “Effectiveness of organizational responses to technological change in the production process.” Organization Science, 3(3): 301-320.
Uzzi, B., and J. Spiro. 2005. “Collaboration and Creativity: The Small World Problem.” American Journal of Sociology, 111 (2): 447–504.
Uzzi, B., S. Mukherjee, M. Stringer, and B. Jones. 2013. “Atypical combinations and scientific impact.” Science, 342(6157): 468-472.
Vakili, K. 2016. “Collaborative Promotion of Technology Standards and the Impact on Innovation, Industry Structure, and Organizational Capabilities: Evidence from Modern Patent Pools.” Organization Science, 27(6): 1504-1524.
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 (6): 1917–1939.
von Hippel, E. 1988. The Sources of Innovation. New York: Oxford University Press.
Wen, W., M. Ceccagnoli, and C. Forman. 2015. “Opening Up Intellectual Property Strategy: Implications for Open Source Software Entry by Start-up Firms.” Management Science, 62 (9): 2668–2691.
Zuckerman, H. 1972. “Age, Aging, and Age Structure in Science, Reprinted in: Robert K. Merton, 1973.” The Sociology of Science.

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

Notes: The ranking on the left indicates the level of impact that the fall of Soviet Union had on the subfield. The higher the subfield's ranking, the more it was impacted 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)

Panel A: Full Sample

Variable	**N**	**Mean**	**Std. Dev.**
Citation-weighted number of publications per year	123,139	5.724	35.193
Simple publication count per year	123,139	0.732	1.299
Number of breakthrough publications (in top 5% cited) per year	123,139	0.043	0.275
Number of breakthrough publications (in top 10% cited) per year	123,139	0.090	0.402
Number of collaborators between 1980 and 1988	123,139	0.541	1.519
Number of unique collaborators between 1980 and 1988	123,139	0.413	0.937

Panel B: Matched Sample

Variable	**N**	**Mean**	**Std. Dev.**
Citation-weighted number of publications per year	81,762	4.678	20.178
Simple publication count per year	81,762	0.690	1.151
Number of breakthrough publications (in top 5% cited) per year	81,762	0.037	0.229
Number of breakthrough publications (in top 10% cited) per year	81,762	0.081	0.346
Number of collaborators between 1980 and 1988	81,762	0.505	1.327
Number of unique collaborators between 1980 and 1988	81,762	0.412	0.920

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

Panel A: Full Sample

Variable	**Specialists**	**Generalists**	**t-test difference**
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*** (p=0.000)
Total number of publications between 1980 and 1988	8.047 (9.413)	8.856 (10.386)	-0.809** (p=0.002)
Total number of breakthrough publications (in top 5% cited) between 1980 and 1988	0.356 (1.116)	0.560 (1.579)	-0.204*** (p=0.000)
Total number of breakthrough publications (in top 10% cited) between 1980 and 1988	0.780 (1.710)	1.102 (2.244)	-0.322*** (p=0.000)
Average number of collaborators between 1980 and 1988	0.703 (1.386)	0.588 (0.729)	0.115*** (p=0.000)
Average number of unique Collaborators between 1980 and 1988	0.446 (0.572)	0.480 (0.517)	-0.034** (p=0.019)

Panel B: Matched Sample

Variable	**U.S. scientists**	**Non-U.S. scientists**	**t-test difference**
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 (p=0.389)
Total number of publications between 1980 and 1988	7.261 (5.395)	7.337 (5.323)	-0.075 (p=0.655)
Total number of breakthrough publications (in top 5% cited) between 1980 and 1988	0.340 (0.910)	0.376 (0.891)	-0.036 (p=0.204)
Total number of breakthrough publications (in top 10% cited) between 1980 and 1988	0.816 (1.537)	0.835 (1.490)	-0.019 (p=0.694)
Average number of collaborators between 1980 and 1988	0.528 (0.772)	0.533 (0.631)	0.018 (p=0.823)
Average number of unique Collaborators between 1980 and 1988	0.410 (0.520)	0.448 (0.478)	0.037** (p=0.017)

Table 4. Differential Creative Output of Specialists Relative to Generalists in Areas Affected by the Collapse of the Soviet Union

Dependent variable	Citation-weighted count of publications	Simple count of publications	Quality-adjusted count of publications	Simple count of Publications
Sample	Full sample	Full sample	Matched sample	Matched sample
Estimation model	Panel Poisson	Panel Poisson	Panel Poisson	Panel Poisson
	(1)	(2)	(3)	(4)
Specialist × SovietRich × AfterIronCurtain (β_1)	0.605*** (0.205)	0.315*** (0.113)	0.777*** (0.253)	0.390*** (0.129)
Specialist × AfterIronCurtain (β_2)	-0.254** (0.103)	-0.078** (0.032)	-0.241*** (0.080)	-0.056 (0.038)
SovietRich × AfterIronCurtain (β_3)	-0.454*** (0.176)	-0.264*** (0.099)	-0.610*** (0.208)	-0.382*** (0.114)
Controls for cumulative publications and non-linear age profile	Yes	Yes	Yes	Yes
Individual and year fixed effects	Yes	Yes	Yes	Yes
No. of observations	113,406	113,512	76,783	76,795
No. of mathematicians	6,132	6,140	4,024	4,024
Chi^2	203.81***	1059.76***	104.60***	663.53***
Log-likelihood	-786985.39	-104275.08	-451570.52	-68821.16

Notes: The data is a panel at the author level based on publication data between 1980 and 2000. The unit of analysis is author-year. All models are Poisson with robust standard errors, clustered at the individual author level. 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. * p < 0.10, ** p < 0.05, ** p < 0.01.

Table 5. Differential Breakthrough Output of Specialists Relative to Generalists in Areas Affected by the Collapse of the Soviet Union

Dependent variable	Count of breakthroughs (publications in top 5% cited)	Count of breakthrough (publications in top 10% cited)	Count of breakthroughs (publications in top 5% cited)	Count of breakthrough (publications in top 10% cited)
Sample	Full sample	Full sample	Matched sample	Matched sample
Estimation model	Panel Poisson	Panel Poisson	Panel Poisson	Panel Poisson
	(1)	(2)	(3)	(4)
Specialist × SovietRich × AfterIronCurtain (β_1)	0.557** (0.257)	0.461** (0.203)	0.685* (0.369)	0.736** (0.292)
Specialist × AfterIronCurtain (β_2)	-0.285*** (0.094)	-0.315*** (0.068)	-0.265** (0.114)	-0.343*** (0.082)
SovietRich × AfterIronCurtain (β_3)	-0.227 (0.224)	-0.171 (0.178)	-0.412 (0.316)	-0.455* (0.252)
Controls for cumulative publications and non-linear age profile	Yes	Yes	Yes	Yes
Individual and year fixed effects	Yes	Yes	Yes	Yes
No. of observations	76,795	77,686	20,641	113,406
No. of mathematicians	4,024	4,076	1,075	6,132
Chi^2	192.90***	338.44***	97.87***	200.33***
Log-likelihood	-10427.39	-19880.25	-6399.016	-12668.48

Notes: The data is a panel at the author level based on publication data between 1980 and 2000. The unit of analysis is author-year. All models are Poisson with robust standard errors, clustered at the individual author level. 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. * $p < 0.10$, ** $p < 0.05$, ** $p < 0.01$.

Table 6. Differential Number of Collaborators of Specialists Relative to Generalists in Areas Affected by the Collapse of the Soviet Union

Dependent variable	Total number of collaborators	Total number of unique collaborators	Total number of collaborators	Total number of unique collaborators
Sample	Full sample	Full sample	Matched sample	Matched sample
Estimation model	Panel Poisson	Panel Poisson	Panel Poisson	Panel Poisson
	(1)	(2)	(3)	(4)
Specialist × SovietRich × AfterIronCurtain (β_1)	0.340** (0.139)	0.428*** (0.157)	0.379** (0.157)	0.457** (0.179)
Specialist × AfterIronCurtain (β_2)	-0.074** (0.037)	-0.110** (0.046)	-0.057 (0.055)	-0.100* (0.054)
SovietRich × AfterIronCurtain (β_3)	-0.270** (0.119)	-0.260** (0.132)	-0.371*** (0.135)	-0.396** (0.155)
Controls for cumulative publications and non-linear age profile	Yes	Yes	Yes	Yes
Individual and year fixed effects	Yes	Yes	Yes	Yes
No. of observations	76,795	96,917	65,986	113,406
No. of mathematicians	4,024	5,243	3,459	6,132
Chi^2	196.63***	122.91***	116.44***	63.56***
Log-likelihood	-71016.25	88857.85	-48285.65	-57694.92

Notes: The data is a panel at the author level based on publication data between 1980 and 2000. The unit of analysis is author-year. All models are Poisson with robust standard errors, clustered at the individual author level. 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. * p < 0.10, ** p < 0.05, ** p < 0.01.

Figure 1. Estimated difference in the quality-adjusted output of specialists versus generalists in the most affected areas of theoretical mathematics, relative to the difference between specialists and generalists in the least affected areas.

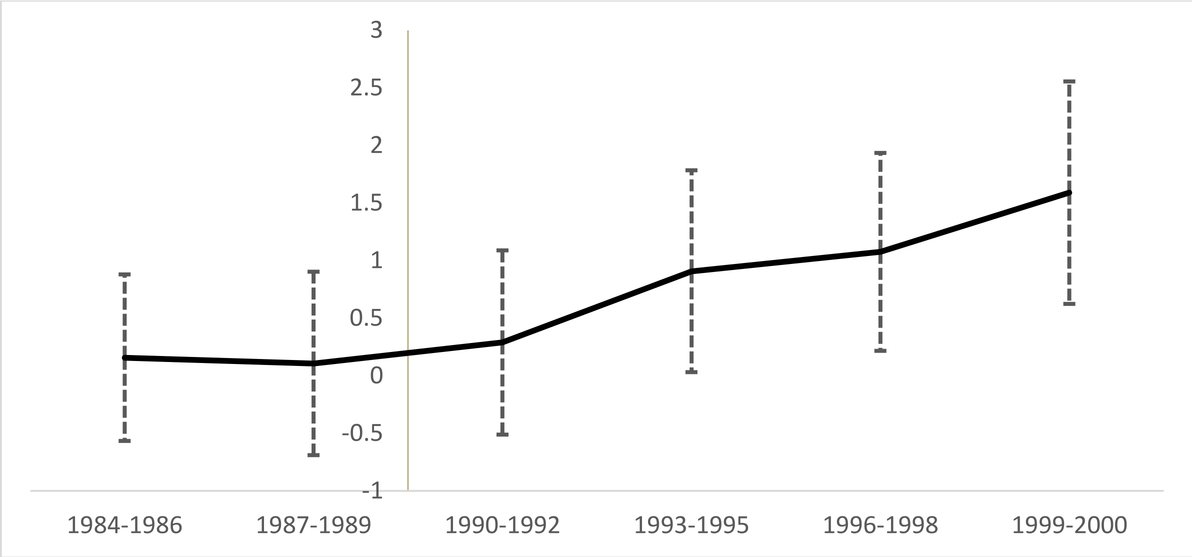

Notes: We base this figure on 10 years of publication data before the collapse of the Soviet Union and 10 years after the collapse. Each point on the graph represents the coefficient value on the covariate *Specialist* × *SovietRich* × *TimePeriod* and thus describes the relative difference in quality-adjusted publication rates between specialists and generalists in areas most affected by the Soviet shock and the same difference in areas least affected. The bars surrounding each point represent the 95% confidence interval. Note that the larger confidence intervals of post-Soviet estimates (relative to our main specification) is 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 1981–1983.

Figure 2. Estimated difference in breakthrough output (top 5% cited publications) of specialists versus generalists in the most affected areas of theoretical mathematics, relative to the difference between specialists and generalists in the least affected areas.

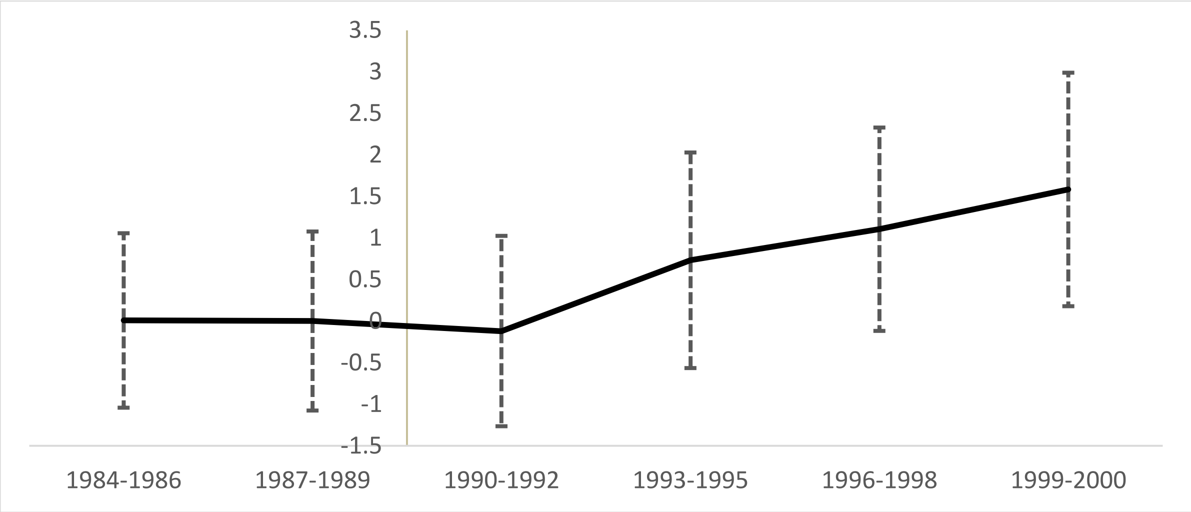

Notes: We base this figure on 10 years of publication data before the collapse of the Soviet Union and 10 years after the collapse. Each point on the graph represents the coefficient value on the covariate *Specialist* × *SovietRich* × *TimePeriod* and thus describes the relative difference in breakthrough output (count of top 5% cited publications) between specialists and generalists in areas most affected by the Soviet shock and the same difference in areas least affected. The bars surrounding each point represent the 95% confidence interval. Note that the larger confidence intervals of post-Soviet estimates (relative to our main specification) is 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 1981–1983.

Figure 3. Estimated difference in breakthrough output (top 10% cited publications) of specialists versus generalists in the most affected areas of theoretical mathematics, relative to the difference between specialists and generalists in the least affected areas.

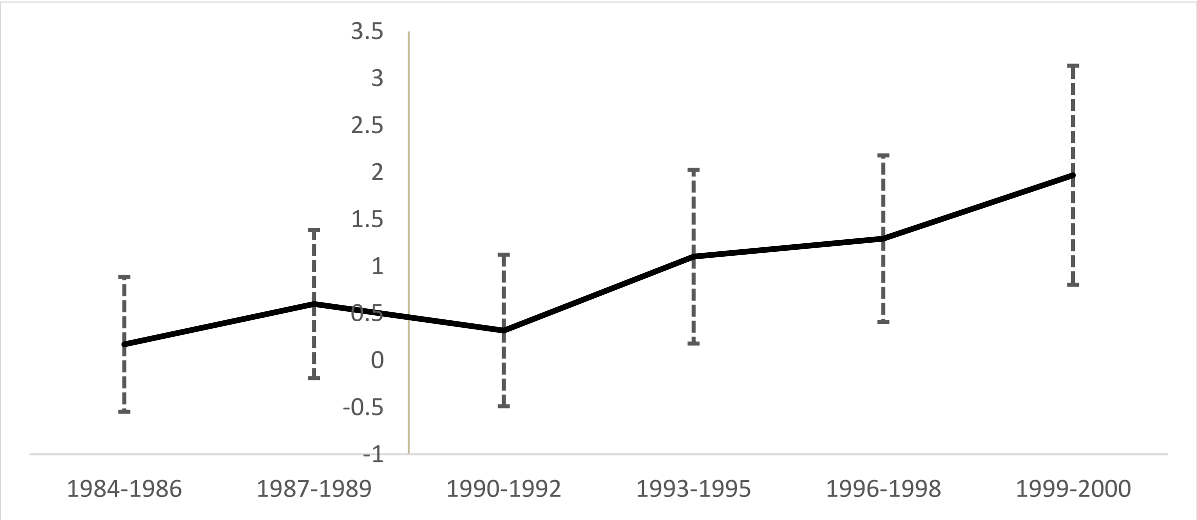

Notes: We base this figure on 10 years of publication data before the collapse of the Soviet Union and 10 years after the collapse. Each point on the graph represents the coefficient value on the covariate *Specialist* × *SovietRich* × *TimePeriod* and thus describes the relative difference in breakthrough output (count of top 10% cited publications) between specialists and generalists in areas most affected by the Soviet shock and the same difference in areas least affected. The bars surrounding each point represent the 95% confidence interval. Note that the larger confidence intervals of post-Soviet estimates (relative to our main specification) is 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 1981–1983.

Figure 4. Estimated difference in collaboration (count of co-authorship instances) of specialists versus generalists in the most affected areas of theoretical mathematics, relative to the difference between specialists and generalists in the least affected areas.

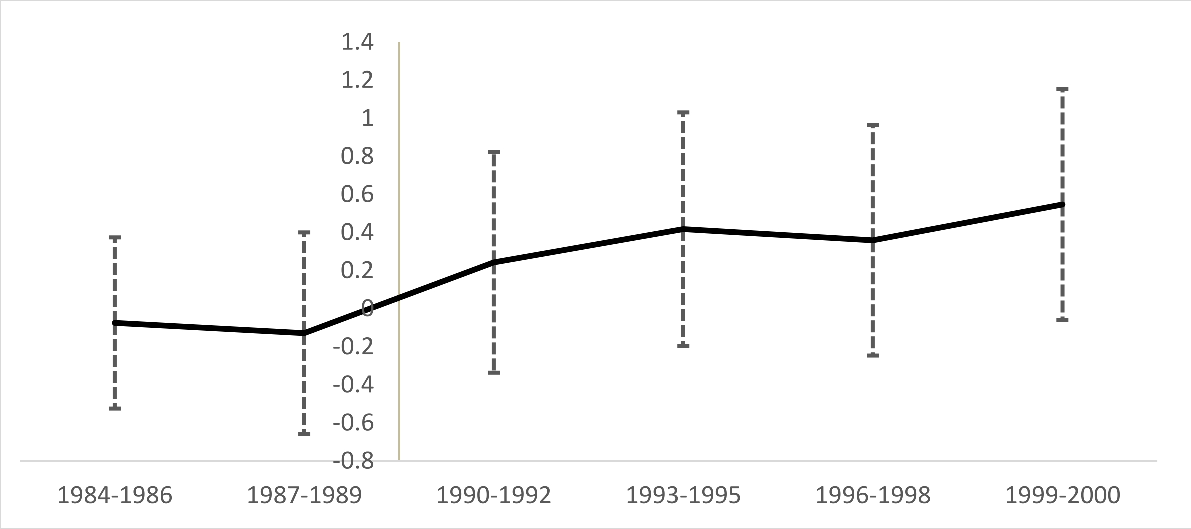

Notes: We base this figure on 10 years of publication data before the collapse of the Soviet Union and 10 years after the collapse. Each point on the graph represents the coefficient value on the covariate *Specialist* × *SovietRich* × *TimePeriod* and thus describes the relative difference in collaboration (count of co-authorship instances) between specialists and generalists in areas most affected by the Soviet shock and the same difference in areas least affected. The bars surrounding each point represent the 95% confidence interval. Note that the larger confidence intervals of post-Soviet estimates (relative to our main specification) is 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 1981–1983.

Figure 5. Estimated difference in collaboration (count of unique collaborators) of specialists versus generalists in the most affected areas of theoretical mathematics, relative to the difference between specialists and generalists in the least affected areas.

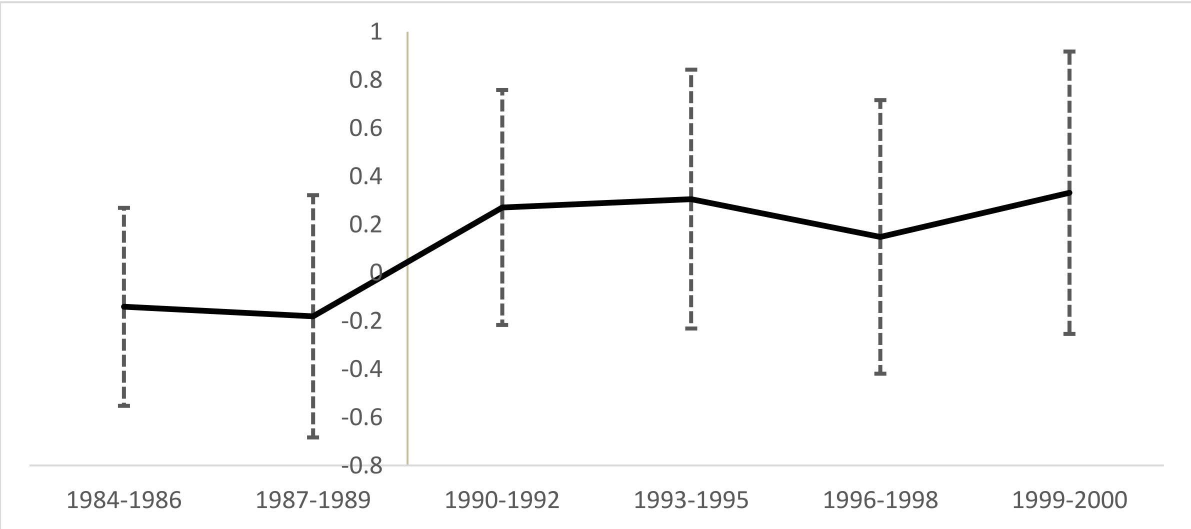

Notes: We base this figure on 10 years of publication data before the collapse of the Soviet Union and ten years after the collapse. Each point on the graph represents the coefficient value on the covariate *Specialist* × *SovietRich* × *TimePeriod* and thus describes the relative difference in collaboration (count of unique collaborators) between specialists and generalists in areas most affected by the Soviet shock and the same difference in areas least affected. The bars surrounding each point represent the 95% confidence interval. Note that the larger confidence intervals of post-Soviet estimates (relative to our main specification) is 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 1981–1983.