# ENTREPRENEURSHIP AND INNOVATION POLICY AND THE ECONOMY:

# INTRODUCTION TO VOLUME 4

Benjamin Jones and Josh Lerner

This volume is the fourth installment in the National Bureau of Economic Research (NBER) *Entrepreneurship and Innovation Policy and the Economy* (EIPE) series. Entrepreneurship and innovation are central forces for long-term prosperity, introducing new products, services, and business models to the economy that raise standards of living and improve human health. Entrepreneurship and innovation also depend critically on government policies and institutions, from the allocation of research grants and tax credits to the regulation of intellectual property to investments in human capital. The annual EIPE conference and associated volume synthesize key research findings about contemporary entrepreneurship and innovation topics. The goal is to distill recent research on critical issues and convey key insights into the policymaking community. The EIPE series also highlights important, open issues that call for future research.

This year's meeting was held in Washington, DC in May 2024, and the six contributions, collected in this volume, address major topics in entrepreneurship and innovation. The first three chapters consider the rise of artificial intelligence. These chapters consider, in turn: the competitive landscape and market structure that may develop as AI advances; the role of intellectual property rules in shaping the advance of AI; and Chinese government investment in the AI sector. Collectively, these three contributions address key dimensions in understanding who will capture value as artificial intelligence advances, and how the nature of this value capture may accelerate or impede technological progress.

The second set of papers include two chapters on research funding models. The first examines the design of Operation Warp Speed (OWS), which delivered effective vaccines in record time, and asks whether and how the OWS model could be applied to other needs. The second considers the design of scientific research grants more broadly, using insights from economics frameworks to analyze existing grant systems and surface key tradeoffs and open research issues. The final chapter of the volume focuses on national security, highlighting the role of the science, technology, engineering and mathematics (STEM) workforce in supporting the US defense sector and examining U.S. workforce needs and immigration policy through this lens. Collectively, the chapters in this volume draw insights from economics and related fields to inform key, contemporary policy questions, while further highlighting critical areas for future research.

In “Old Moats for New Models: Openness, Control, and Competition in Generative AI,” Pierre Azoulay, Josh Krieger, and Abhishek Nagaraj examine the future evolution of artificial intelligence, focusing on how its industry structure is likely to evolve. To do so, the authors draw lessons from earlier episodes involving breakthrough innovations as well as the framework that economists have developed to understand these events.

In particular, the authors highlight the importance of two considerations that are likely to shape the structure of the artificial intelligence industry. The first of these is appropriability: the ability of firms to obtain proprietary rights to their discoveries. The second of these is the concept of complementary assets, or the extent to which firms also need other critical assets to successfully exploit artificial intelligence. The authors argue that the second consideration—assets such as computing power, safety protocols, and access to massive quantities of non-public training data—will be the critical “moat” that allows incumbent firms to dominate the market. If this future scenario is to be avoided, the authors argue, it is likely to depend on a “rogue” technology giant choosing an open architecture as a way to fight their rivals.

“Intellectual Property and Creative Machines,” by Gaétan de Rassenfosse, Adam Jaffe, and Joel Waldfogel, focuses in more depth on the first of the barriers that Azoulay and co-authors identified: intellectual property. In particular, they focus on how generative artificial intelligence will interact with copyright protection. The authors highlight the pressures the copyright system is under today: the extent to which copyright covers artificial intelligence models is extremely uncertain, yet the economic stakes are huge. Software developers are potentially engaging in copyright infringement on a massive scale by using unlicensed data to train generative AI models, but obtaining the relevant rights to these training data is likely to be prohibitively expensive for all developers except those with the deepest pockets. Meanwhile, the copyright system was not designed for the generative AI world, raising questions about the extent to which the output of these models is protected.

The authors highlight that despite these concerns, there are costs to acting either too late or too soon. In favor of moving quickly on rulemaking, uncertainty may deter investments or reward firms that infringe on others in the hopes that they will be forgiven for their past transgressions during any legislative process. On the other hand, given the rapid evolution of artificial intelligence models, premature efforts to formulate rules that turn out to be problematic will be difficult to change. While the authors do not have clear answers, they suggest a set of questions that should be helpful to policymakers navigating these treacherous waters.

Martin Beraja, Wenwei Peng, David Yang and Noam Yuchtman's essay, "Government as Venture Capitalists in AI," the third and final work in this set of papers, examines the development of artificial intelligence in the second great hub for its development, China. This nation, seeking global leadership in AI, has set up venture funds sponsored by both local and national governments. These funds seek to harness the power of venture investing—its emphasis on intensive screening and monitoring, staged financing, and high-powered incentives—to advance state aims.

The authors empirically contrast China's government funds with traditional and private venture capital funds. They highlight several striking patterns. In particular, the government funds are far more spread out geographically than the private funds, and willing to invest in seemingly weaker firms. After the investments, however, the government fund-backed firms enjoys superior growth and attract considerable private investment. The authors suggest that, at least in the Chinese context, government venture funds may be well positioned to overcome the information problems that typically represent a major challenge to successful venture investment.

The next two chapters consider research funding models. In "Can Operation Warp Speed Serve as a Model for Accelerating Innovations Beyond Covid Vaccines?," Arielle D'Souza, Kendall Hoyt, Christopher Snyder, and Alec Stapp elucidate the key features of OWS and consider the conditions under which the OWS approach can generalize to other settings. The chapter begins by surveying "innovation missions" in U.S. history, where the U.S. government made unusually large and rapid public investments in research and development with specific technological goals. This analysis helps put OWS in broader context – with the Apollo program, the Manhattan Project, and others – and reveals factors that make such public investments possible. Key factors include national importance, time sensitivity, the need for coordination, the insufficiency of the commercial market, and well-defined technological objectives.

The chapter then considers key design features of OWS per se, providing a detailed history of the effort's goals and methods. One key feature was the portfolio approach to investment, where OWS simultaneously took "many shots on goal" to increase the chances that at least one effective vaccine became available in a short period of time. The authors discuss the importance of trying "longshots" and emphasize that unlikely-but-uncorrelated pathways can substantially increase the odds of overall success. Indeed, the mRNA platforms that proved especially effective were initially seen as especially unlikely to work, yet they enabled OWS to deliver vaccines in record time. The chapter further identifies and examines numerous other important features of OWS, including the use of both push and pull funding mechanisms, the standing up of manufacturing facilities in

parallel with the research pathways, the roles of coordination and leadership, and the authority to disrupt standard protocols and regulatory constraints in pursuit of speed. The chapter further positions these features in terms of the enormous health and economic costs of pandemic and the social benefits that could be realized through an effective vaccine. Finally, the chapter considers how design features of OWS may or may not apply to other contexts. Applications to Alzheimer's disease and climate change are discussed. One feature of OWS that appears broadly relevant is the "many shots on goal" approach, while other features appear less applicable.

In "Designing Scientific Grants," Christoph Carnehl, Marco Ottaviani, and Justus Preusser consider a major component of innovation policy – grants to support scientific research. The root design challenge is to allocate limited available research funding across a potentially large set of researchers and their projects. Difficulties ensue because funders and researchers have somewhat different interests – the researcher primarily seeks funding for themselves while the funder wants to choose only the best projects. The parties also have different sets of information – in particular, the researcher has private information about the quality of their ideas and their effort in execution. To address these challenges, the authors provide a framework that highlights key features of the funding systems, including the application, evaluation, and monitoring phases. At each phase, the authors engage careful consideration of the incentive effects and costs of different design approaches – and the unintended consequences they may impose. To make progress, the authors apply established insights from information economics and mechanism design to the specific tradeoffs that emerge through the different phases of the grant funding process.

One important finding concerns proportional allocation rules, where agencies (such as the National Institutes of Health or the European Research Council) work to equalize application success rates across fields. At first blush this seems like a reasonable and fair approach, but it can have unintended consequences. In particular, researchers are less likely to apply for grants that they are less likely to receive. As such, fields where evaluation of projects is more accurate tend to see less applicants, because researchers with less meritorious applications are less likely to bother to apply, knowing they will be rejected. But this can create the unintended consequence of shrinking budgets for these fields: namely, since the agency is working to equalize applicant success rates across fields, fewer applicants in a field means fewer projects will be funded in that field. Funders then end up penalizing fields where evaluation is more accurate. The authors detail this conceptual issue and further demonstrate it using European Research Council data, showing that funding through this proportional process has resulted in funds being shifted away from the life sciences and toward the social science and humanities, where the value of research

prospects is less clear. The surprising implication is that better evaluation in a given field, when embedded in these funding systems, ends up reducing funding for that field. More broadly, this chapter presents a large array of such insights across different phases of the funding process. In taking cutting-edge theoretical tools to the grant-funding context, the chapter provides novel and deep perspectives on grant funding design, identifying key tradeoffs as well as important new areas for research.

The last chapter turns to human capital. In "Meeting U.S. Defense Science and Engineering Workforce Needs: A Progress Report," Amy Nice focuses on immigration policy in light of the U.S. defense sector's substantial demand for STEM workers. Previous economics research has shown the central roles of foreign-born workers in U.S. science, invention, and entrepreneurship. This chapter focuses on the role of foreign-born in supporting U.S. national security. Indeed, staying technologically ahead of other nations is critical to supporting national security, so that STEM workers appear essential to defense and global leadership. Further, while 14% of the U.S. population is foreign-born, 37% of U.S. STEM workers involved in defense projects are foreign born, indicating how much the talent base for U.S. national security draws from global talent pools. This chapter begins by reviewing the role of foreign-born talent in supporting U.S. national security and identifies this topic as an important hole in economics research on immigration, which has been more concerned with non-defense applications. The chapter then reviews immigration policy in light of these defense considerations.

A key contribution of the chapter is to encompass the manifold set of visa categories for U.S. immigration, define their objectives and constraints, and then consider numerous specific options to advance STEM immigration in support of national security, whether through regulatory rules or legislative statutes. Deeply informed by detailed institutional knowledge, the chapter provides a compendium of potential, tangible policy innovations. Finally, the paper considers empirical approaches, both reviewing recent work and identifying fundamental holes in available data and analysis, which can inform the understanding of STEM workers and their national security implications. Altogether, this paper extends analysis of foreign-born STEM workers in the U.S. from their role in U.S. economic prosperity to U.S. national security, provides an organized and clear synthesis of policy options, and defines future research needs.

The six contributions in this volume all engage with high-stakes, contemporary issues in innovation and entrepreneurship policy. The evolution of artificial intelligence, the design of research funding systems, and STEM workforce needs in U.S. national security are all at the leading edge of policy discussions. The chapter authors have synthesized the research literature and applied core economics frameworks to address these critical, contemporary

issues. Each contribution provides rich insights, highlighting policy options and deepening foundations for researchers and policymakers who seek to better understand and advance the U.S. innovation and entrepreneurship system and meet core national priorities.