

Selecting Innovation Funding Mechanisms: A Framework for Policymakers

Matthew Clancy, Coefficient Giving, United States of America

Matthew Esche, Institute for Progress, United States of America

Siddhartha Haria, Market Shaping Accelerator, United States of America

Claire McMahon, Harvard University, United States of America

Christopher M. Snyder, Dartmouth College and NBER, United States of America

Caleb Watney, Institute for Progress, United States of America

July 2026

Abstract

The U.S. federal government spends roughly \$150 billion a year on research and development through mechanisms ranging from grants and contracts to prizes, advance market commitments, and tax credits. These funding mechanisms differ markedly in risk allocation, decision authority, and incentive power, yet policymakers have little systematic guidance for choosing among them. We develop a framework that asks the policymaker three initial questions about the contracting environment: Can they articulate the work to be done? Can they specify what would constitute a successful solution? Can they identify the most capable performers to contract with? The answers considerably narrow the set of viable mechanisms, enabling a more tractable choice of the efficient mechanism based on further considerations from the innovation-economics literature.

JEL Codes: O38, O32, L52

Keywords: innovation policy, research and development, public funding, verifiability, asymmetric information, grants, prizes, advance market commitments

Acknowledgments: The authors thank Patrick Collard and Maxwell Tabarrok for research assistance, Sarrin Chethik for contributions to this paper and the companion tool, Jana Gallus and Heidi Williams for guidance on the project from early stages, and participants at the 2026 NBER Entrepreneurship and Innovation Policy and the Economy Conference for helpful comments. The paper benefitted from the suggestions of editors Ben Jones and Josh Lerner. All remaining errors are our own. The Institute for Progress gratefully acknowledges support from Coefficient Giving and the Smith Richardson Foundation. The Market Shaping Accelerator gratefully acknowledges support from the Sijbrandij Foundation and the Center for Incentive Design. Snyder gratefully acknowledges support from Coefficient Giving and Irving Institute for Energy and Society at Dartmouth. The “Atlas of Innovation” (Institute for Progress and Market Shaping Accelerator 2026) is an online tool that is a companion for this paper, available at atlasofinnovation.org.

I. Introduction

In the mid-2000s, NASA awarded a contract to Kistler Aerospace to resupply the International Space Station. The contract was subsequently withdrawn amid legal and procurement challenges, prompting NASA to reconsider both the structure and competitiveness of its approach (French 2013). NASA ultimately replaced it with the Commercial Orbital Transportation Services (COTS) program, an open competition paying fixed-price awards upon successful completion of predefined development milestones (U.S. Government Accountability Office 2009). This shift in funding mechanism enabled new entrants such as SpaceX to develop capabilities that have since become central to the commercial space sector.

The COTS experience suggests that the design of the funding program can matter as much as the level of funding for the success of an innovation program. Different mechanisms materially affect which innovators participate and which capabilities emerge. The U.S. federal government spends roughly \$150 billion annually on research and development (R&D), shaping the nation's scientific frontier and national competitiveness; yet it has no systematic framework to help policymakers choose the innovation-funding mechanism best suited to their context.

One might hope that designing such a framework would be a straightforward task: if a single mechanism were to dominate all the others, the only task would be to identify that mechanism and recommend it. However, examples can be offered of mechanisms that work well in one setting but poorly in other, even ostensibly similar, settings. The Energy Independence and Security Act of 2007 directed U.S. fuel suppliers to blend increasing volumes of biofuels into the fuel supply. The policy applied the same instrument to two technologies, with sharply divergent results. Corn ethanol scaled steadily and met its statutory volume requirements through 2022. Cellulosic ethanol, derived from the fibrous rather than starchy parts of plants, fell far short of its statutory schedule, forcing EPA to revise the required volumes downward repeatedly. The divergence reflects the technological state of the two industries when the mandate was adopted (Clancy and Moschini 2017). Corn ethanol relied on a mature production process; the main remaining challenge was deployment. Guaranteed demand was therefore sufficient to draw capacity into the market. Cellulosic ethanol, by contrast, still required scientific and engineering advances before production could become scalable and commercially viable.

The federal government uses more than a dozen disparate mechanisms to fund innovation, including research grants, procurement contracts, milestone payments, prizes, advance market

commitments, tax credits, loan guarantees, and agencies created to direct a portfolio of projects (ARPAs, short for Advanced Research Projects Agencies). These mechanisms entail markedly different allocations of risk, decision authority, and incentive power, making it difficult for policymakers to choose among them. In the absence of systematic guidance, it is natural for a policymaker to default to mechanisms that the agency has always used. The agency's "business as usual" might be more of a historical accident rather than optimized for current conditions and perforce excludes some of the more novel market-shaping mechanisms. The opposite inclination—favoring mechanisms just because of their novelty—is also problematic, possibly implementing ill-fitting policies that ignore the agency's accumulated wisdom.

An idealized approach to selecting the optimal mechanism would evaluate the benefits and costs of each and select the one delivering the highest net present value. This approach is likely to be impractical here. Undertaking an exhaustive analysis of over a dozen program designs would likely be too time-consuming for a policymaker even if the relevant information were available. In an innovation context, replete with "known unknowns" and "unknown unknowns," there is likely too much uncertainty to undertake such a quantitative exercise. The policymaker may have little idea of the feasible technological pathways, let alone the underlying distribution of innovator capabilities (to assess the potential for adverse selection) or effort costs (to assess the potential for moral hazard) that would be required for the computation of expected costs and probabilities of success across competing mechanisms.

We propose a tractable framework for mechanism selection that begins by characterizing the contractual environment in which the policymaker operates. Three simple questions go a long way toward assessing the contractual constraints facing the policymaker: Is it possible to articulate the problem area performers should work on? Is it possible to specify what would constitute a successful solution? Is it possible to identify the most capable innovators to contract with?

The answers to these questions can quickly narrow down the choice set to a workable number of viable mechanisms that can be analyzed in more detail. In environments that allow relatively complete contracting, our framework highlights the opportunity of leveraging contractual richness to exert more control over the direction and intensity of innovative effort. In other environments, the policymaker may have to delegate more of these decisions to the innovator holding superior local information.

Having delineated viable mechanisms and ordered them in terms of contractual richness, our framework next has the policymaker consider a list of frictions from the innovation-economics literature to identify the most efficient choice within that set. These include financing constraints on performers, interdependence among technical approaches, and R&D spillovers across innovation stages; other familiar frictions, including adverse selection, moral hazard, and tech maturity, are already embedded in our analysis of the three core questions. The framework identifies the most efficient mechanisms in the viable set as those best able to overcome the frictions that are relevant to the particular innovation environment. The last stage takes into account logistical and political considerations that might restrict the use of certain mechanisms in the specific context.

We have adapted this analytical foundation into the “Atlas of Innovation” (Institute for Progress and Market Shaping Accelerator 2026), a companion online tool depicted in Figure 1 that operationalizes the framework by guiding policymakers through progressively detailed questions to arrive at mechanisms that may fit their specific innovation challenge.




Fig. 1. Screen Capture of Companion Online Tool, “Atlas of Innovation.” The user navigates over land and sea to a destination funding mechanism by answering a series of questions about the innovation challenge they have in mind. The landing page for the funding mechanism provides extensive detail on the funding mechanism including definitions, examples, pros and cons, and relevant economic literature. The tool cross-references alternative mechanisms that may also be suitable.

The remainder of this paper proceeds as follows. Section II reviews the relevant literature on innovation policy, information economics, and mechanism design. Section III develops the principal-agent model that underpins our framework and explains how the degree of contractual incompleteness in the innovation setting can limit the set of viable funding mechanisms. Section IV presents our framework in detail, walking through the three core questions noted above, and mapping them to specific funding mechanisms. Section V concludes by discussing policy implications and directions for future research.

II. Literature Review

The case for public innovation funding rests on well-established economic principles. Arrow (1962) showed that the public-good characteristics of knowledge—non-rivalry and partial non-excludability—lead private firms to underinvest in R&D. Innovators capture only a fraction of the social value they create (Nordhaus 2004), and this appropriability problem is compounded when innovations generate unpriced externalities like vaccine herd immunity or climate benefits (Kremer and Glennerster 2004; Jaffe, Newell, and Stavins 2005). These market failures justify intervention, but agency theory demonstrates that the form of intervention matters: adverse selection and moral hazard constrain what policymakers can feasibly and efficiently contract (Akerlof 1970, Arrow 1963, Holmström 1979, Pauly 1968).

To address these market failures in innovation, policymakers have developed a diverse menu of financing tools, with well-documented trade-offs. Scholars have analyzed the utility of grants for early-stage exploration, finding that while traditional peer-review systems can be risk-averse and biased against novel ideas (Azoulay, Graff Zivin, and Manso 2011; Li 2017), grants remain highly effective for financially constrained startups by serving as a certification of quality that unlocks private capital (Howell 2017). The ARPA model combines mission-oriented programs with active management by empowered program directors, enabling breakthrough innovation through portfolio shaping and coordination (Azoulay et al. 2019, Goldstein and Narayanamurti 2018), yet ARPAs can be difficult to replicate institutionally due to rare program manager talent. “Pull” mechanisms, which condition payments on verifiable research outputs, can resolve the moral hazard problem and encourage competition (Brunt, Lerner, and Nicholas 2012; Kremer, Levin, and Snyder 2020), but require success criteria to be specified in advance. Fiscal incentives like R&D tax credits effectively raise aggregate R&D spending (Bloom, Griffith, and Van Reenen 2002; Becker 2015; Dechezleprêtre et al. 2023) but lack the precision to direct innovation toward specific goals.

Several practitioner playbooks have cataloged the breadth of these funding mechanisms (Sellick, Solder, and Roberts 2018; Yu and Yu no date; Renaissance Philanthropy no date), highlighting the benefits, drawbacks, and implementation details of each approach, and academic literature has provided rigorous theoretical comparisons and toolkits (Wright 1983; Bloom, Van Reenen, and Williams 2019; Azoulay and Li 2021; Armitage, Bakhtian, and Jaffe 2024). In other policy domains, researchers have organized policy instrument choice around features of the

environment that policymakers can directly observe, such as whether finance is scarce or infrastructure is lacking, rather than unobservable theoretical primitives (Hausmann, Rodrik, and Velasco 2008). Despite this wealth of analysis, a systematic framework for selecting from the full set of innovation funding mechanisms is missing.

To address this gap, we propose a tractable set of questions based on the information available to the policymaker; specifically, what the policymaker knows about the problem, solution, and performer. This approach builds on core insights from incomplete-contracting theory pioneered by Grossman and Hart (1986): when complete contracts are impractical, the organization may have to be structured to provide substitute investment incentives. Harris and Raviv (1979) and Spier (1992) are early studies of contract design in the presence of monitoring and contracting costs. Most relevant to our work, Lerner and Malmendier (2010) examine the implications of incomplete contracts in the specific context of innovation. Myers, Lanahan, and Johnson (2026) argue that high contractibility is one of the key factors determining the suitability of projects for the Small Business Innovation Research program they study.

III. Conceptual Model

In this section, we introduce a conceptual model that will clarify the objectives of our framework and build a vocabulary to facilitate the discussion.

Consider the principal-agent model drawn schematically in Figure 2. We take the perspective of the principal — a policymaker who seeks to use public funding to support innovation to further some social goal. Goals can take on many forms, such as the generation of new scientific ideas (such as a deeper understanding of the biological processes of aging), the demonstration of technical feasibility (such as a working prototype of a fusion energy device), or the large-scale deployment of a product or service (such as a low-cost vaccine manufactured and distributed at scale). The agents in the model are called “performers,” which can be individuals, firms, research labs, or other organizations. Performers undertake the effort that, when successful, leads to an innovation outcome. Stipulate that the policymaker’s social goal is worthy and that public funding is warranted, presuming that the commercial market will not generate adequate innovation incentives without public support. The policymaker chooses among funding mechanisms to maximize the present value of the benefits of innovation net of program expenditures.




Fig. 2. Structure of the Conceptual Model. The principal (policymaker) offers the agents (performers) a funding mechanism, which incentivizes their investment of innovation effort in pursuit of an innovation outcome.

Since our framework focuses on choosing the type of funding mechanism rather than the amount of funding, we will abstract from mechanism sizing, supposing that the policymaker has a roughly suitable budget required for the scale of the intended innovation project. For other work that focuses on determining the optimal size of an innovation-funding mechanism, see, for example, Snyder, Hoyt, and Gouglas (2023).

The expected benefit from innovation equals the product of two factors: the benefit from achieving an innovation goal and the probability of success. The second factor distinguishes innovation policy from standard procurement: by definition, an innovation has never been achieved before, making failure far more likely than in standard procurement of established products. The more advanced the innovation is beyond the current technology frontier, typically the lower the probability of success. The probability of success depends on how many performers

attempt the innovation, the capabilities of participating performers, the effort they exert, and the direction of that effort—all of which the policymaker's choice of mechanism can affect.

The chosen mechanism sets the terms of the optimization problem facing performers, which determines their innovation strategies. Each potential performer decides whether to participate at all and, if so, how much effort to devote to an approach. The participation decision hinges on whether the expected rewards, including the mechanism's payment and any reputational or property-rights benefits of success, exceed the costs of time, money, and effort. Those performers who decide to participate direct their effort toward the activities and outcomes that the mechanism rewards and choose effort levels that optimize expected costs against expected benefits. The mechanism's payment terms thus determine who attempts an innovation, how hard they work, and where their efforts are directed.

The policymaker can manage performers' innovative effort in two ways. One approach contracts directly on effort inputs, while the other contracts on an output target and leaves performers to choose how best to reach it, including their effort level, technological approach, and other features of the path to the target. The first approach funds work in specified problem areas. The second funds performers after achieving a specified solution target. The second approach rewards success, not attempts; it pays for outputs, not inputs. By having performers bear more of the risk of failure, the second approach has the potential to generate higher-powered incentives.

When the policymaker does not manage the innovative effort through targeting of the problem or solution—paying neither for specific types of attempts nor outcomes—policymakers may delegate to the performer to decide which attempts to undertake and outcomes to target. Policymakers instead can select performers to fund based on performer characteristics or capabilities.

Figure 3 illustrates some of the consequences of each targeting strategy. Problem-targeting defines a boundary in attempt space (first row); solution-targeting defines a target in outcome space (second row). In the third row, neither targeting strategy is possible; and so the policymaker is left to delegate direction entirely to performers.




Fig. 3. Direction of Innovation Effort under Various Targeting Strategies. In each panel, the large white oval represents the space of possible innovation attempts and the white rounded rectangle represents the space of possible outcomes. Arrows map attempts to outcomes. Funded attempts are indicated by black circles, funded outcomes by black rectangles, and the true social optimum by the star. Panels A and B show cases arising under problem-targeting, in which the policymaker chooses a mechanism that contracts on attempts within a specified area (hatched region in attempt space). The area is well-calibrated to the social optimum in A and mistargeted in B. Panels C and D show cases arising under solution-targeting, in which the policymaker only funds performers once their outcomes fall within a specified target (hatched region in outcome space). The target encompasses the social optimum in both panels but is so broad in D that the majority of funded performers miss the optimum by a wide margin. Panels E and F illustrate delegation, in which the policymaker specifies neither a problem area nor solution target, but instead has chosen funded performers through some selection process. Performers choose their own direction, which may align with the social optimum (panel E) or not (panel F).

Whether the mechanism is conditioned on the problem or solution, getting the specification right is critical. The policymaker may hope for panel A or C. Those panels illustrate cases in which the policymaker can specify the problem area or target outcome well, leading to the outcome that is closely aligned with the social optimum (star) with little wastage on misdirected attempts. If the policymaker lacks the information or contractual language to specify the problem area or solutions target well, however, the mechanism may result in poorly directed innovation effort. In panel B, the policymaker narrowly targets a problem area, but the wrong one. In panel D, the policymaker sets such a broad target that most outcomes stray far from the social optimum. When it is difficult for the policymaker to specify the problem area or solution target well, delegating direction to better-informed performers may be more attractive. If there is little divergence between performers' and social preferences, as in panel E, their undirected choice may produce a desirable outcome. As panel F illustrates, delegation is not a panacea: performers may choose directions far from the social optimum due to misaligned private incentives or uncertainty about the optimal path.

The key advantage of solution-targeting over problem-targeting is that it can often generate stronger incentives, which can help select capable performers and induce them to exert concerted effort toward the desired outcome. The drawback is that solution-targeting likely requires more careful and complete contracting. A vague or incomplete contract could leave scope for the policymaker to expropriate the returns from the performer's innovation effort. For example, the policymaker could renege on the commitment to reward innovation and save the expenditure by claiming that the performer fell short of the vague solution target; a court may have difficulty deciding otherwise. Anticipating their effort may not receive its full reward, performers would naturally underinvest in innovative effort, an instance of the so-called hold-up problem (Grout 1984, Grossman and Hart 1986). Vaguely specifying "work to improve pandemic preparedness" may suffice for a low-powered, problem-targeted incentive to conduct research in a broad area, but to condition success or failure on the outcome of investment, more precise terms such as "develop a universal vaccine that is at least 75% effective against two coronavirus genera" might be required, offering less "wiggle room" out of the commitment to reward success. The set of mechanisms that can be used in a particular setting depends on the completeness of contracts the policymaker can write there.

Discussions of contractual frictions often center on asymmetric-information problems, chiefly adverse selection and moral hazard. However, uncertainty need not be asymmetric to impede contracting. Innovations may emerge only after decades of work, under conditions very different from those prevailing when the contract is written. Policymakers and performers may therefore share deep uncertainty about a long list of known unknowns that would be difficult to specify in advance, as well as unknown unknowns that cannot be specified at all.

Funding mechanisms incentivize innovation efforts in the face of these information frictions. A mechanism specifies who is paid, when payment occurs, and what conditions trigger payment. These design choices affect which performers enter, how much effort they exert, which approaches they pursue, and how much risk they bear. Project grants, researcher-based grants, contracts, prizes, advance market commitments, and tax credits differ because they allocate authority over problems, solutions, and performers in different ways. For example, project-based grants require the funder to assess the problem, proposed solution, and performer; researcher-based grants specify the performer but leave problems and approaches more open; and pull mechanisms specify success criteria while using competition to reveal capable performers.

Table 1 catalogs the funding mechanisms analyzed in this paper and offers examples of them in use. The definitions reflect common usage, but the boundaries between mechanisms are not always clear-cut, and implementation often varies in ways that no single definition can fully capture. The categories are meant to provide a working vocabulary for the analysis, not an exhaustive or mutually exclusive typology. The companion online tool, the “Atlas of Innovation,” provides a more complete treatment of each mechanism.

Table 1

Catalog of Innovation Funding Mechanisms

Funding mechanism	Definition	Example
Advance market commitment (AMC)	Fund set aside to provide top-up payment for a yet-to-be-invented product meeting specified technical requirements	GAVI Pneumococcal AMC: \$1.5 billion commitment to subsidize future purchases of vaccines ^a
Coordinated research program	Institutional funding framework designed to accelerate high-risk research through active program management	Defense Advanced Research Projects Agency: empowered program directors manage portfolios toward defined breakthroughs, including ARPANET and GPS ^b
Field building	Strengthening entire disciplines, sectors, or professional ecosystems by training new talent, creating resources, and building legitimacy	Semiconductor Research Corporation: industry, academic, and government consortium building shared workforce pipelines and technical infrastructure ^c
Research joint venture	Collaboration between government agency and external research institution, sharing resources and expertise to tackle scientific challenges or develop breakthrough technologies	Argonne National Laboratory–Constellation Energy: shared government research infrastructure and private nuclear assets to develop carbon-free power technologies ^d
Intramural science	Research conducted by in-house teams, often in funder’s own facilities	National Institutes of Health Intramural Research Program: in-house scientists identified HIV and enabled the first effective treatments ^e
Loan instruments	Reducing innovators’ borrowing costs by direct lending on favorable terms, subsidizing interest rates on private loans, or guaranteeing repayment to lenders if borrowers default	Department of Energy Loan Programs Office: \$465 million guaranteed loan to Tesla for electric vehicle manufacturing scale-up ^f
Milestone payments	Staged funding disbursements rewarding incremental progress toward an innovation target	NASA COTS: milestone-based contracts enabled less-capitalized companies, including SpaceX, to develop private spacecraft ^g
Per-unit subsidy, per-unit tax credit	Predetermined payment for each unit of an innovative product produced or purchased, creating incentives that scale with output or adoption	Renewable Electricity Production Tax Credit: per-kilowatt-hour credit for qualifying renewable energy generation ^h
Prize, absolute	Lump-sum reward for innovation meeting technical target	Ansari XPRIZE: \$10 million award for first private reusable spacecraft to reach 100km altitude twice in two weeks ⁱ
Prize, relative	Lump-sum reward for innovation outperforming peers ^j	DARPA Grand Challenge: \$2 million award to fastest autonomous vehicle completing a defined desert course ^j

Procurement contract	Agreement to purchase innovative product from specified supplier	UK Department of Health: fixed annual fee to pharmaceutical companies for access to effective antibiotics ^k
R&D contract	Agreement between funder and performer (often a company) to conduct R&D with clearly specified timeline, budget, and deliverables	NASA–Lockheed Martin: \$247.5 million contract to design and test the X-59 quiet supersonic aircraft ^l
R&D tax credit	Companies reduce their tax bill by claiming credits for eligible R&D expenses	US Research & Experimentation Tax Credit: broad credit reducing after-tax cost of qualifying R&D across industries ^m
Research grant, project-based	Funding to individuals or organizations to support research on a specific research question with a defined method or goal, often provided upfront or as cost reimbursements	National Science Foundation Digital Libraries Initiative: grant funding PageRank research at Stanford, which became the foundation for Google ⁿ
Research grant, researcher-based	Funding to individuals or organizations to cover research costs for a broad research agenda, often provided upfront or as cost reimbursements	Howard Hughes Medical Institute Investigator Program: seven-year renewable awards for self-directed research ^o

Sources: All examples are drawn from the “Atlas of Innovation” (Institute for Progress and Market Shaping Accelerator 2026). More detail on each example can be found in the following sources: ^agavi.org/investing-gavi/innovative-financing/pneumococcal-amc, ^barpa.mil, ^csrc.org, ^danl.gov/partnerships/collaborative-partnerships, ^eirp.nih.gov, ^fenergy.gov/EDF, ^gU.S. Government Accountability Office (2009), ^hepa.gov/lmop/renewable-electricity-production-tax-credit-information, ⁱxprize.org/competitions/ansari, ^jarpa.mil/about/innovation-timeline/grand-challenge, ^kcommonslibrary.parliament.uk/netflix-for-antimicrobials-the-antimicrobial-products-subscription-model/, ^lnasa.gov/news-release/nasa-awards-contract-to-build-quieter-supersonic-aircraft/, ^mirs.gov/businesses/research-credit, ⁿnsf.gov/impacts/google, ^ohhmi.org/programs/investigators.

Note: Authors’ catalog of public funding mechanisms commonly used to support innovation, listed in alphabetical order. Each program can take a variety of forms, but for brevity we focus on typical forms of these agreements.

IV. Framework for Mechanism Selection

The NASA COTS example showed that, for a given innovation context, some mechanisms may perform substantially better than others. Conversely, the ethanol example showed that the performance of a given mechanism can be sharply different even in contexts that are ostensibly quite similar. The optimal funding mechanism depends on potentially subtle features of the innovation setting. With 15 mechanisms and scores of factors to weigh for each, policymakers are understandably daunted by the task of how best to spend their innovation budgets.

This section develops a tractable framework for policymakers to select among funding mechanisms.

We call a mechanism “viable” if the policymaker can write an enforceable contract and has adequate information to make well-informed decisions about its clauses and parties. Our framework assesses contractual completeness — and thus viability — through three simple questions: Is it possible to articulate the problem area the policymaker wants performers to work in? Can the policymaker specify a verifiable solution target? Can the policymaker identify capable performers ex ante? When all three answers are yes, relatively complete contracts are possible and a wide range of mechanisms are viable. When one or more answers are no, the viable set narrows and the policymaker must resort to mechanisms that delegate decision authority, induce self-selection, or reward outcomes without specifying the path to achieving them.

As illustrated in Figure 4, the framework proceeds in stages. The first stage asks the three core questions determining which mechanisms are viable. The answers determine which mechanisms are even viable. The next stage brings in additional considerations from the innovation-economics literature to narrow down the viable set of mechanisms further and assess their relative efficiency. The last stage looks for political and logistical constraints that might rule out some mechanisms from being feasible to implement.




Fig. 4. Three-Stage Framework for Selecting Funding Mechanism. The first stage asks three core questions about the extent of contractual completeness in the setting, determining the set of viable mechanisms. The second stage brings in further economic considerations to narrow the set of mechanisms within the viable set. The last stage considers practical details that might preclude certain mechanisms.

A. Contractibility Questions

Our framework starts with three core questions to assess the completeness of contracts in the setting. A funding mechanism will not even be viable if it requires information that the policymaker does not have or cannot reliably elicit from performers. If the policymaker does have sufficiently accurate and verifiable information, mechanisms that condition funding, selection, or rewards on that information can strengthen effort incentives and identify more capable performers, overcoming moral hazard and adverse selection. Thus, full use of available information tends to make mechanisms more efficient.

The three contractability questions do not provide a complete picture of the intricacies of a funding mechanism decision, but they are a straightforward way to narrow the possible menu of options and elicit what a policymaker knows in the relevant context.

We discuss the three core questions about contractual completeness in turn.

Articulate Problem?

The first question asks whether the policymaker can articulate the problem area they want performers to work on. Policymakers can have general objectives such as “advance the field of quantum computing” or “improve understanding of aging.” These may be important goals but do not lend themselves to contracting on deliverables or selecting in advance the endeavors most likely to succeed.

When policymakers cannot clearly articulate the problem, they can resort to mechanisms that let performers direct their own innovation efforts. Capable performers (if such exist) may know more about which approaches are promising and tractable and when the activities should shift in response to scientific and technological advances in their subfield.

When the policymaker cannot articulate the problem area, they cannot write contracts that specify deliverables tightly or condition payment on outcomes defined in advance. This narrows the viable set of funding mechanisms to those leaving performers discretion to define research directions and priorities: R&D tax credits, field building, researcher-based grants, and intramural science. If in addition to being unable to articulate the problem area, the policymaker cannot identify capable performers *ex ante*, the viable set narrows further, leaving only R&D tax credits and field building.

When policymakers can articulate problems precisely, on the other hand, they can contract on which questions are in scope, and allow for contracting on deliverables, enabling project-based grants and research contracts. If contracting is even more complete, allowing them to specify outputs and verify when they have been achieved, this opens the door to prizes, AMCs, and procurement.

Specify Solution?

The second question asks whether the policymaker can specify the solution they want performers to provide. In other words, can contracts be conditioned on verifiable output targets? Having to specify a solution is a more demanding contracting problem than articulating a problem area to work on. A policymaker may be able to precisely articulate the problem that needs solving but have no basis for describing the features a solution must possess. For example, a policymaker may know that artificial intelligence systems pose safety risks without being able to specify what a verifiably safe system would look like or how to test whether it has been achieved.

When a solution cannot be specified, this precludes performance-based mechanisms that define criteria for success in advance. Which remaining mechanisms are viable depends on what else can be specified in a contract. In the space of incomplete contracts where the policymaker can articulate the problem but not the form the solution should take, there is high technical uncertainty, multiple approaches may be viable, and the optimal end product is unknown. This favors exploratory mechanisms that fund progress in a general problem domain rather than toward a specific solution, such as directed R&D contracts, ARPA-style program management, and research joint ventures. These mechanisms allow researchers to pivot as new information emerges and to pursue approaches the policymaker could not have anticipated. These mechanisms fund legitimate attempts to solve an innovation problem irrespective of whether that attempt succeeds.

On the other hand, when contracting is more complete, allowing the policymaker to specify the features a successful solution must possess, contracts can condition payment on outcomes. Target product profiles, performance specifications, and minimum efficacy thresholds can be included in the contract to render success verifiable, even when no solution yet exists.

One virtue of contracting on success rather than attempts is that it can specify the endpoint while leaving the technological path open, allowing innovators who possess relevant domain

expertise the discretion over how to reach the endpoint. The policymaker can remain agnostic as to which firm or which technology is most suitable, allowing competition among firms and new technologies to flourish.

Mechanisms that pay for results include absolute prizes, AMCs, milestone payments, and per-unit subsidies. Each depends on defining success up front and conditioning payment on achieving it.

Identify Capable Teams?

The third question asks whether there is any way to identify capable teams *ex ante*, or whether that information will only emerge later as the mechanism operates. Two complementary pieces of information are required: which technical approaches are most promising, and who is best positioned to deliver them. The policymaker does not have to possess this information personally and can instead obtain it through auctions, requests for proposals, and expert panels. The question is whether such processes yield reliable signals.

The policymaker must decide whether to select teams *ex ante*, to sign a bilateral contract before innovation begins, or to use open mechanisms that pay upon completion, which let potential performers decide whether to participate. Direct selection is only valuable when two conditions hold. First, a pool of capable performers already exists. Without trained researchers or capable firms, field building may need to precede project funding. Second, a policymaker must be able to identify the most capable among them. This is possible either when performers are roughly equivalent (any selection then suffices) or when observable signals such as track records, proposals, or peer review correlate meaningfully with capability. When these conditions are met, grants and contracts that rely on direct selection become more attractive.

When these conditions fail, policymakers cannot reliably distinguish the most capable teams. Innovators hold private information about their own capabilities and costs that they cannot credibly communicate, and the policymaker lacks the technical knowledge to evaluate such claims. Further, policymakers may not be able to monitor costs or effort exerted by the R&D performer. This favors a contracting mechanism that induces self-selection and pays contingent on delivery rather than bilateral contracts agreed with specific performers *ex ante*. Open mechanisms with this feature include prizes and AMCs. Because these mechanisms allocate technological risk to innovators,

they attract performers with capability, cost advantages, or large non-monetary payouts. Performers self-select based on their own assessments of expected costs and their probability of success, information that the policymaker cannot access directly.

When policymakers can identify the most promising performers, direct contracting and grant-making become viable. The policymaker has sufficient information, through track records, auctions, or a request for proposals process, to contract specific performers without relying on open competition. This avoids having to compensate firms for the risk that competition entails and reduces duplicative effort across teams pursuing the same goal. It also enables coordination across performers in ways that open competition does not. Mechanisms for this scenario include grants, contracts for R&D, procurement contracts, and coordinated research programs.

Synthesizing Answers to the Three Core Questions

Table 2 uses the answers to the three core questions to delineate viable funding mechanisms. Higher-numbered rows require more contractability but open up more mechanisms as viable. Field building and R&D tax credits are viable in any event, while procurement contracts, project-based research grants, and loan instruments only become viable when the policymaker can answer all three questions affirmatively.

Table 2

Set of Viable Mechanisms Determined by Answer to Three Core Questions

Row number	Answer to core question			Viable mechanisms (efficient starred)
	Articulate problem?	Specify solution?	Identify performer?	
(1)	No	Not applicable	No	*Field building *R&D tax credit
(2)	No	Not applicable	Yes	Field building R&D tax credit *Researcher-based research grant *Intramural science†
(3)	Yes	No	No	Field building R&D tax credit *Relative prize
(4)	Yes	No	Yes	Field building R&D tax credit Researcher-based research grant Intramural science *Coordinated research program *Research joint venture *R&D contract
(5)	Yes	Yes	No	Field building R&D tax credit Relative prize *Absolute prize *Milestone payments *Advance market commitment *Per-unit subsidy

(6)	Yes	Yes	Yes	Field building R&D tax credit Researcher-based research grant Intramural science Coordinated research program Research joint venture R&D contract Relative prize Absolute prize Milestone payments Advance market commitment Per-unit subsidy *Procurement contract *Project-based research grant *Loan instruments
-----	-----	-----	-----	---

Note: Table lists viable funding mechanisms for each possible configuration of answers to core questions. Higher-numbered rows reflect richer contractual environments, so more “yes” answers leading to these environments unlock more viable mechanisms. Mechanisms exploiting the richer information available to the funder in higher numbered rows are designated “information efficient” with asterisks; this designation is a starting point, refined by later stages of the framework bringing in additional considerations. Rows (1) and (2) implicitly assume that articulating a problem is less contractually demanding than specifying a solution, so answering “no” to the first question precludes an independent answer to the second question, hence the entry “not applicable.” †The version of intramural science envisioned in row (2) is its open-ended form, where in-house researchers set their own agendas within a broad mission area. In practice, intramural science is sometimes implemented for logistical considerations such as security requirements or long-term institutional needs, considerations apart from the contractual questions dealt with later in the framework.

For tractability, the answers to the questions are treated as yes/no judgments; but in practice these are matters of degree. The table also implicitly assumes that articulating a problem area is less contractually demanding than specifying a target solution.

Starred mechanisms within a row carry a presumption of greater efficiency: they can be conditioned on more information and thus provide higher powered incentives and more direction to innovation. This presumption is a starting point rather than a conclusion. As the selection framework proceeds to incorporate further economic factors and practical considerations, non-starred mechanisms in a row may emerge as being relatively more suitable in context. That additional analysis, which we turn to next, will also help the policymaker select among the starred mechanisms appearing together in a row.

B. Further Economic Considerations

The previous subsection focused on the main role of the three core questions: determining the completeness of contracts and thus which mechanisms are viable and efficient in the relevant setting. In fact, while we frame the three core questions around the single dimension of contractability, they capture deeper economic forces than that framing suggests. Those three simple questions operationalize a range of important concepts from information economics and innovation policy.

For example, whether a policymaker can articulate the problem and specify the solution depends on the technological maturity of the domain. Whether performers can be identified ex ante implicitly reflects adverse selection: when capability signals are strong, direct selection is efficient; when signals are weak, mechanisms must induce self-selection through contract terms. These information constraints also determine whether outcome-based contracts can mitigate moral hazard: when solutions cannot be specified in verifiable terms, policymakers must accept weaker incentives and fund exploration rather than predetermined endpoints.

Therefore, in addition to narrowing the set of mechanisms to a smaller viable set, the three core questions already make progress toward incorporating key economic concepts (technological maturity, adverse selection, moral hazard) in a manner that is understandable to many policymakers. Of course, three simple questions cannot exhaust the range of economic considerations. In this subsection, we discuss further considerations from innovation economics—financing constraints, coordination needs, R&D spillovers—that arise in specific innovation contexts.

Financing Constraints

Mechanisms that defer payments—pull mechanisms that condition payment on successful delivery or push mechanisms structured as reimbursements—require performers to finance their own effort until those payments are received. This financing burden grows with the time horizon between investment and payment. When projects require years of development before deliverables arrive, only performers with access to patient capital can participate.

The timing delay makes self-financed projects expensive because performers typically face higher costs of capital than government funders. Private firms, startups, and university labs must borrow or raise equity at market interest rates that exceed the government's cost of borrowing (DiMasi, Grabowski, and Hansen 2016; Gormsen and Huber 2025). As the time between initial investment and payment lengthens, this interest rate differential compounds, making pull mechanisms substantially more expensive than push mechanisms that use government funds from the outset. Capable but cash-constrained performers may be excluded even when they hold the most promising approaches.

This creates a trade-off between effort incentives and participation. Outcome-contingent payment strengthens the link between effort and reward, increasing effort intensity among participants. But when financing constraints bind, the participation effect can dominate: fewer performers attempt the problem, and the excluded performers may have been the most likely to succeed. In such cases, pull mechanisms can reduce the overall probability of innovation success by reducing the number of participants even as they increase effort per participant.

Several mechanisms can help mitigate the financing burden. Milestone payments break a long project into stages, providing intermediate financing and payment while also generating verifiable signals that can attract private capital. Loan instruments (including direct government loans, loan guarantees, and interest rate subsidies) provide capital access that preserves an incentive structure through repayment obligations while directly leveraging the government's lower cost of capital. Grants can eliminate financing burdens by providing payment upfront, and although this weakens outcome incentives, repeated grantmaker–grantee interactions can mitigate these concerns through reputational incentives and expectations of future funding. The appropriate choice depends on how severely financing constraints restrict the pool of capable performers.

Interdependence in Approaches

Technical approaches to an innovation goal may be interdependent in two ways: they may share common failure modes, or they may benefit from coordination due to complementarities in capabilities.

When approaches share common failure modes, their outcomes are correlated. If one approach fails due to a shared scientific assumption or regulatory barrier, others relying on the

same foundation are likely to fail as well. In this case, funding many parallel attempts provides less diversification than it may appear. The policymaker who believes they are hedging risk by supporting multiple performers may instead be concentrating resources on a single underlying bet.

Addressing this correlation problem requires greater policymaker discretion in selecting performers and approaches, rather than relying on open mechanisms like prizes or AMCs. When the innovation task is relatively straightforward and duplication of effort is unnecessary, a targeted single contract may be appropriate. When multiple approaches are needed, policymakers can either concentrate resources on a small number of genuinely distinct technical paths or use active program management, such as ARPA-style programs or curated grant portfolios, to steer performers toward complementary rather than redundant approaches.

When approaches benefit from coordination among performers, mechanisms that reward individual achievement impose costs by discouraging information sharing. If progress on one subproblem would accelerate progress on others, or if performers hold complementary capabilities, competitive mechanisms leave coordination gains unrealized. ARPA-style programs address this through active program management, shared infrastructure, and structured information exchange. Joint ventures formalize collaboration between parties with complementary assets, specifying terms for intellectual property and data sharing. Intramural research eliminates organizational boundaries, enabling tight coordination but sacrificing the diversity and independent judgment that external performers may provide. Alternatively, prizes can be structured to reward intermediate outputs that reveal progress and facilitate information sharing across competitors before committing to full-scale efforts.

R&D Spillovers

Innovation often occurs in stages—initial invention followed by scale-up and commercialization—potentially performed by different actors. If the domain is amenable to forms of intellectual property that protect the returns of early-stage innovators' investments, this can both strengthen early-stage performers' investment incentives and facilitate the flow of information to later innovators through licensing contracts (Sampat and Williams 2019). In such settings, suitable funding mechanisms do not have to protect early innovators' investments from expropriation or

micromanage the flow of information along intermediate stages of the innovation pipeline but can instead focus on incentivizing final outcomes.

When early-stage contributions generate spillovers, knowledge spreads to the whole field rather than staying with the originator. Mechanisms that only acknowledge final outcomes will inadequately incentivize the crucial first steps. In these cases, mechanisms that explicitly reward each stage have advantages. Milestone payments can address this by providing intermediate rewards for breakthroughs that enable subsequent innovation, even when those breakthroughs cannot be protected or sold.

C. Practical Considerations

The first two stages of the selection framework discussed so far generate a small set of potentially efficient funding mechanisms. The final stage of the framework, discussed in this subsection, checks whether one of a list of practical considerations would preclude the otherwise efficient mechanism's use, requiring the policymaker to choose a different one.

Legal and Political Constraints

Agencies are bound to operate within legal limits, where statutory mandates and procurement regulations may rule out mechanisms that would otherwise be attractive. For example, until the America COMPETES Reauthorization Act of 2010 granted broad prize authority, many federal agencies lacked statutory authority to run prize competitions. Additionally, appropriations rules generally constrain agencies from disbursing funds beyond the fiscal year availability that Congress sets at time of appropriation, which complicates mechanisms such as AMCs that require committing money years before any payout. Implementing an AMC may require an outside fund holder, or the policymaker may need to substitute a mechanism that pays within the budget window.

Political considerations matter as well. Budget scoring conventions can make some mechanisms harder to adopt than others, regardless of their economic merits, and mechanisms differ in how easily they are sustained once adopted. Before a pay-for-success mechanism pays out, the activity it induces can be hard to observe, making the program difficult to defend to

stakeholders who expect visible progress, even when the mechanism is working exactly as designed.

Institutional Fit

Innovation settings differ in which type of performer is best positioned to do the work given their environment or institutional fit. In some cases, different aspects of the performer's institutional environment determine feasibility of a mechanism — whether due to security, trust, institutional continuity, or commercial incentives.

Constraints on performer environments are often relevant for security-sensitive work. Some innovation problems involve information or infrastructure that cannot be widely shared, such as classified defense and intelligence data, sensitive personal information such as health or census records, or technologies like nuclear weapons whose details must be tightly held. Security requirements of this kind constrain who can perform the work and under what oversight.

These requirements push the policymaker toward using mechanisms such as intramural science, government–research lab joint ventures, or R&D contracts with closely vetted performers. For example, the national laboratories' stewardship of the nuclear stockpile relies on trusted institutions with specialized capabilities, continuity, and security requirements. Beyond security, in-house researchers are also suitable for long-term programs requiring institutional stability, and settings where in-house expertise protects analysis from outside interests, as when the Congressional Budget Office scores legislation or the Food and Drug Administration reviews drug applications.

In other cases, institutional environments shape performers' incentives and behavior. Academic researchers are well-suited to create knowledge, particularly through long-term fundamental research, and to disseminate it and train others. However, they are less well-suited to commercializing their research and are likely less responsive to commercial rewards. Private-sector firms, by contrast, have strengths in bringing products to market, scaling production, and responding to market incentives, but leveraging these strengths requires them to capture some of the value their innovations generate.

Mechanisms tend to work better when compatible with the capabilities and incentives of the performers that an innovation setting calls for. A pull mechanism conditioned on market success

may do little to motivate university scientists doing basic research, just as a researcher-based grant is often a poor instrument for inducing a firm to scale manufacturing.

Unitary Versus Atomistic Consumers

When the innovation outcome is a product for final consumers, the nature of those consumers affects which mechanisms are best-suited. If the government is the main purchaser and the policymaker is funding the development of a product for its own use, procurement contracts may have unique advantages. By pre-committing to purchase, the government provides a strong pull incentive, and because the contract is written by the end-user, its requirements can be tailored to actual needs.

If a commercial market exists, the policymaker need not bear the full cost of research and development; private demand will carry part of the burden. Loan instruments and per-unit subsidies make efficient use of public funds in this setting, easing financing frictions or bridging the gap between private and social value while leaving the rest to the market. AMCs sit between these cases, specifying a future per-unit subsidy and potentially a guaranteed purchase to establish a market floor.

Duplicative Funding

The framework so far treats the policymaker as if they were the only funder, but innovation funding can be a crowded field: other agencies, philanthropies, private investors, and foreign governments may already be supporting work toward the same goal. Mechanism choice may therefore account for how a new program interacts with existing funding.

Overlapping funding programs can act as substitutes when they target the same margin of activity. A prize for achieving a goal adds little if existing grants already fund ample effort toward that same goal, and public funding of any kind may be unnecessary if commercial returns are lucrative enough to attract private capital on their own.

Alternatively, funding programs can act as complements when they target different margins or stages. For example, grants can fund the early exploration that no performer would risk on a prize alone, while a prize or AMC pulls the resulting discoveries through to a finished outcome.

The question for the policymaker is whether existing funding programs already induce the activity that the new mechanism would pay for.

Other funders may include foreign governments, whose programs compete for mobile research activity. R&D tax credits reward research conducted in the sponsoring jurisdiction, so their effectiveness depends on what other jurisdictions offer. A credit less generous than foreign alternatives may attract little activity, while a more generous one may simply pay firms to relocate research they would have conducted anyway.

V. Conclusion

We propose a framework for selecting among innovation funding mechanisms that puts the policymaker and what they know about the completeness of contracts front and center, determining which mechanisms are even viable. The framework assesses the completeness of contracts in the relevant setting by asking the policymaker three questions about their information environment: Can they precisely articulate the problem domain? Can they precisely specify the solution? Can they identify the most capable performers upfront? These questions determine which mechanism classes are viable and which decisions may be best delegated to better-informed performers in the market.

Within the viable set, additional factors help narrow the choice. Financing constraints favor mechanisms that reduce the time between investment and payment (milestones, loan instruments, or upfront grants) over pure pull mechanisms that require performers to self-finance. Interdependence among approaches favors coordinated mechanisms when information sharing is essential for success and argues against naive diversification when approaches share common failure modes. When innovators generating knowledge early along the innovation path have difficulty preventing spillovers to other performers in the same or later stages, mechanisms that compensate early innovators directly can promote investment in each innovation stage.

The framework provides only a partial guide to innovation policy. It assumes that one or another market failure prevents the commercial market from delivering the right amount of innovation without some external funding. The policymaker is assumed to have correctly determined that public funding of innovation is warranted after carefully weighing costs and benefits. As such, the framework can inform how to pursue a goal once chosen, but not which goals deserve public funds in the first place. Our framework is silent on the appropriate level of

public funding to achieve an innovation goal. Our framework takes for granted that public funding is the appropriate policy intervention; alternatives such as strengthening intellectual property protections, streamlining immigration for high-skilled workers, or reforming regulation may be more appropriate in some settings.

Our framework focuses on ex ante decisionmaking and abstracts from program experimentation and evaluation. A variety of mechanisms could be piloted in different contexts to gain a deeper understanding about what works when. Even more could be learned if such pilots could be organized as randomized controlled trials, although running scores of large innovation experiments may be prohibitively expensive in many cases. Much could be learned from ex post evaluations of programs that have been run.

Future work could test the framework's predictions empirically, examining whether mechanism choices aligned with information environments produce better innovation outcomes than misaligned choices. Dynamic extensions could address how mechanism choice should evolve as information is revealed over the course of a research program. And richer models could incorporate multiple funders, strategic performer interaction, and the endogenous development of policymaker information over time.

Selecting among innovation funding mechanisms is a far more complex decision than any framework can fully capture. Our aim has been to reduce the problem from a daunting menu of mechanisms and considerations into a tractable three-question viability structure. We hope this framework serves as a starting point for policymakers navigating real funding decisions and for economists building toward a more complete theory of innovation policy. Because funding mechanisms are one of the primary ways innovation policy is administered, the choice of the right mechanism can turn an aspirational goal into successful innovation.

References

- Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84 (3): 488–500.
- Armitage, Sarah, Noël Bakhtian, and Adam B. Jaffe. 2024. "Innovation Market Failures and the Design of New Climate Policy Instruments." *Environmental and Energy Policy and the Economy* 5 (1): 4–48.
- Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by the Universities-National Bureau Committee for Economic Research, Committee on Economic Growth of the Social Science Research Council. Princeton University Press.
- Arrow, Kenneth. 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review* 53 (5): 941–73.
- Azoulay, Pierre, Erica Fuchs, Anna P. Goldstein, and Michael Kearney. 2019. "Funding Breakthrough Research: Promises and Challenges of the 'ARPA Model.'" *Innovation Policy and the Economy* 19 (1): 69–96.
- Azoulay, Pierre, Joshua S. Graff Zivin, and Gustavo Manso. 2011. "Incentives and Creativity: Evidence from the Academic Life Sciences." *RAND Journal of Economics* 42 (3): 527–54.
- Azoulay, Pierre and Danielle Li. 2021. "Scientific Grant Funding." In *Innovation and Public Policy*, edited by Austan Goolsbee and Benjamin Jones. University of Chicago Press.
- Becker, Bettina. 2015. "Public R&D Policies and Private R&D Investment: A Survey of the Empirical Evidence." *Journal of Economic Surveys* 29 (5): 917–42.
- Bloom, Nicholas, Rachel Griffith, and John Van Reenen. 2002. "Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997." *Journal of Public Economics* 85 (1): 1–31.
- Bloom, Nicholas, John Van Reenen, and Heidi Williams. 2019. "A Toolkit of Policies to Promote Innovation." *Journal of Economic Perspectives* 33 (3): 163–184.
- Brunt, Liam, Josh Lerner, and Tom Nicholas. 2012. "Inducement Prizes and Innovation." *Journal of Industrial Economics* 60 (4): 657–96.
- Clancy, Matthew, and GianCarlo Moschini. 2017. "Mandates and the Incentive for Environmental Innovation." *American Journal of Agricultural Economics* 100 (1): 198–219.
- Dechezleprêtre, Antoine, Elias Einiö, Ralf Martin, Kieu-Trang Nguyen, and John Van Reenen. 2023. "Do Tax Incentives Increase Firm Innovation? An RD Design for R&D, Patents, and Spillovers." *American Economic Journal: Economic Policy* 15 (4): 486–521.

- DiMasi, Joseph A., Henry G. Grabowski, and Ronald W. Hansen. 2016. "Innovation in the Pharmaceutical Industry: New Estimates of R&D Costs." *Journal of Health Economics* 47: 20–33.
- French, George D. 2013. "George D. French Oral History Transcript." Interview by Rebecca Hackler. NASA Johnson Space Center Oral History Project, Commercial Crew & Cargo Program Office, May 1. <https://www.nasa.gov/wp-content/uploads/2025/06/frenchgd-5-1-13.pdf?emrc=bb24e4>.
- Goldstein, Anna P. and Venkatesh Narayanamurti. 2018. "Simultaneous Pursuit of Discovery and Invention in the US Department of Energy." *Research Policy* 47 (8): 1505–12.
- Gormsen, Niels Joachim and Kilian Huber. 2025. "Corporate Discount Rates." *American Economic Review* 115 (6): 2001–49.
- Grossman, Sanford J. and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy* 94 (4): 691–719.
- Grout, Paul A. 1984. "Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach." *Econometrica* 52 (2): 449–60.
- Harris, Milton and Artur Raviv. 1979. "Optimal Incentive Contracts with Imperfect Information." *Journal of Economic Theory* 20 (2): 231–59.
- Hausmann, Ricardo, Dani Rodrik, and Andrés Velasco. 2008. "Growth Diagnostics." In *The Washington Consensus Reconsidered: Towards a New Global Governance*, edited by Narcís Serra and Joseph E. Stiglitz. Oxford University Press.
- Holmström, Bengt. 1979. "Moral Hazard and Observability." *Bell Journal of Economics* 10 (1): 74–91.
- Howell, Sabrina T. 2017. "Financing Innovation: Evidence from R&D Grants." *American Economic Review* 107 (4): 1136–64.
- Institute for Progress and Market Shaping Accelerator. 2026. "Atlas of Innovation." Online tool available at atlasofinnovation.org.
- Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins. 2005. "A Tale of Two Market Failures: Technology and Environmental Policy." *Ecological Economics* 54 (2): 164–74.
- Kremer, Michael and Rachel Glennerster. 2004. *Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases*. Princeton University Press.
- Kremer, Michael, Jonathan Levin, and Christopher M. Snyder. 2020. "Advance Market Commitments: Insights from Theory and Experience." *American Economic Association Papers and Proceedings* 110: 269–73.

- Lerner, Josh and Ulrike Malmendier. 2010. “Contractibility and the Design of Research Agreements.” *American Economic Review* 100 (1): 214–46.
- Li, Danielle. 2017. “Expertise versus Bias in Evaluation: Evidence from the NIH.” *American Economic Journal: Applied Economics* 9 (2): 60–92.
- Myers, Kyle R., Lauren Lanahan, and Evan E. Johnson. 2026. “Small Business Innovation Applied to National Needs,” *Entrepreneurship and Innovation Policy and the Economy* 5: 97–132.
- Nordhaus, William D. 2004. “Schumpeterian Profits in the American Economy: Theory and Measurement.” NBER working paper no. 10433.
- Pauly, Mark V. 1968. “The Economics of Moral Hazard: Comment.” *American Economic Review* 58 (3): 531–37.
- Renaissance Philanthropy. “Playbooks.” Online tool available at <https://www.renaissancephilanthropy.org/playbooks>.
- Sampat, Bhaven and Heidi L. Williams. 2019. “How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome.” *American Economic Review* 109 (1): 203–36.
- Sellick, Vicki, Amy Solder, and Isobel Roberts. 2018. “Funding Innovation: A Practice Guide.” London: Nesta Foundation. Online article available at <https://media.nesta.org.uk/documents/Funding-Innovation-Nov-18.pdf>.
- Snyder, Christopher M., Kendall Hoyt, and Dimitrios Gouglas. 2023. “An Optimal Mechanism to Fund the Development of Vaccines Against Emerging Epidemics.” *Journal of Health Economics* 91: 102795.
- Spier, Kathryn E. 1992. “Incomplete Contracts and Signalling.” *RAND Journal of Economics* 23 (3): 432–43.
- U. S. Government Accountability Office. 2009. “NASA: Commercial Partners Are Making Progress, but Face Aggressive Schedules to Demonstrate Critical Space Station Cargo Transport Capabilities.” Report to Congressional Addressees GAO-09-618. Washington, DC. Online article available at <https://www.gao.gov/assets/gao-09-618.pdf>.
- Wright, Brian D. 1983. “The Economics of Invention Incentives: Patents, Prizes, and Research Contracts.” *American Economic Review* 73 (4): 691–707.
- Yu, Kelvin and Anson Yu. “Levers for Progress.” Online tool available at <https://leversforprogress.com>.