Conference draft, comments welcome.

Selecting Innovation Funding Mechanisms: A Framework for Policymakers

By MATTHEW CLANCY, MATTHEW ESCHE, SIDDHARTHA HARIA, CLAIRE MCMAHON, CHRISTOPHER SNYDER, AND CALEB WATNEY

The U.S. federal government spends roughly $150 billion annually on research and development using a variety of distinct funding mechanisms, from grants and contracts to prizes, advance market commitments, and tax credits. Each mechanism allocates risk, decision rights, and incentives differently, yet there is no systematic framework for matching mechanisms to innovation challenges. We develop an information-based framework that reframes mechanism selection as a classification problem organized around three questions: Can the policymaker articulate the problem to be solved? Can the policymaker characterize what a successful solution looks like in verifiable terms? Can the policymaker identify the most capable performers ex ante? The answers determine which dimensions of an innovation project the policymaker can contract on and which must be delegated to better-informed innovators, narrowing the list of mechanisms to a smaller set of viable and information-efficient mechanisms. The framework offers tractable guidance for policymakers to apply what they know about the observable features of their innovation environment to identify a feasible set of funding mechanisms.

* Clancy: Coefficient Giving (email: matt.clancy@coefficientgiving.org). Esche: Institute for Progress (email: matthew@ifp.org). Haria: University of Chicago (email: siddharthaharia@uchicago.edu). McMahon: Harvard University (email: clairemcmahon@g.harvard.edu). Snyder: Dartmouth College (email: christopher.m.snyder@dartmouth.edu). Watney: Institute for Progress (email: caleb@ifp.org). Acknowledgments. The authors are grateful to Patrick Collard and Maxwell Tabarrok for early research assistance that informed this paper, to Sarrin Chethik for thoughtful comments on this paper and his intellectual collaboration on the companion tool, and to Heidi Williams, whose guidance and insights have shaped the project from its earliest stages. All remaining errors are our own. The Institute for Progress gratefully acknowledges support from the Smith Richardson Foundation, the Market Shaping Accelerator gratefully acknowledges support from the Sijbrandij Foundation, and Snyder gratefully acknowledges support from the Sloan Foundation under G-2025-25183. An interactive companion to this paper is available at atlasofinnovation.org. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

1. 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 and Hackler 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 (GAO 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 a public funding mechanism can determine the success or failure of an innovation program. Different mechanisms materially affect who participates 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; but it has no systematic framework to help policymakers choose between different funding mechanisms for different innovation challenges.

Economic theory emphasizes that contract design matters under uncertainty and asymmetric information; but in practice, mechanism choice is often relegated to an administrative detail rather than treated as a core policy decision. Standard economic approaches to this problem invoke concepts like moral hazard, adverse selection, and risk. These are analytically useful concepts, but they require information that policymakers rarely possess ex ante.

The federal government uses more than a dozen distinct ways to financially support innovation, including traditional research grants, procurement contracts, milestone-based agreements, prizes, advance market commitments (AMCs), tax credits, loan guarantees, and ARPA-style portfolio programs. Each mechanism allocates risk, decision rights, and incentives differently, and most were introduced to address specific historical or institutional needs. Yet there is little guidance for policymakers on how to match these

tools to the structure of the innovation problem at hand. In the absence of such guidance, policymakers tend to rely on familiar instruments embedded in existing budgeting, oversight, and administrative systems.

In principle, one might hope to select funding mechanisms by comparing their expected costs and benefits against each other, holding the total budget fixed. In practice, this approach is rarely tractable: innovation outcomes are highly uncertain; technological risks are difficult to quantify ex ante; and innovators often hold key private information on feasibility, effort, and alternative approaches. These features make it nearly impossible to estimate the magnitude of informational asymmetries or to assign reliable probabilities to success across competing mechanisms. As a result, policymakers cannot realistically perform a full cost-benefit analysis over the space of available funding tools.

We propose a tractable approach that reframes mechanism choice as a classification problem. Instead of ranking mechanisms by expected value, policymakers can ask which features of the innovation environment favor one class of mechanisms over another. This can identify which funding tools are appropriate given what the policymaker and potential innovators know, what can be specified or verified ex ante, and where uncertainty is most acute.

Reframing the choice of funding mechanism as a classification problem helps — but without paring down the space of possible choices, the classification itself remains almost unworkably complex. Funding mechanisms differ along a variety of dimensions: how they allocate technological and financial risk, reward effort versus outcomes, accommodate learning and adaptation, and mitigate moral hazard or adverse selection. The particular innovation problem also has its own unique features, varying by technological distance from the frontier, degree of uncertainty about feasibility, need for coordination across actors, and extent to which performance can be observed *ex post*.

These dimensions are analytically correct, but they only become useful once the policymaker is able to diagnose the latent information problems in their policy context. In practice, policymakers are rarely well-positioned to perform this diagnosis. Each of these considerations has intuitive appeal and is emphasized in different strands of the economics

literature, but, taken together, they generate a high-dimensional decision problem with little traction on narrowing the set of plausible tools. Even more fundamentally, the policymaker may not know which of these economic constructs is a practical concern in their policy context.

Our approach reduces this complexity by replacing theoretical constructs with observable, actionable questions. We focus on the fundamental information constraints that recur across innovation settings and organize our framework around three core questions.

- *Articulate the problem:* Can the policymaker precisely articulate the problem to be solved — that is, does a policymaker have sufficient information to specify which objectives should be pursued? Or, conversely, does the policymaker know only generally in what direction they would like to push innovation?
- *Characterize the solution:* Can the policymaker precisely characterize the solution to that problem by defining observable and verifiable outcomes to condition funding on? Or is the policymaker reliant on others to identify success?
- *Identify the performers:* Can the policymaker identify the most capable performers *ex ante*? Or will capability only gradually be revealed through experimentation and performance?

Together, these questions identify which dimensions of an innovation project the policymaker has sufficient information to condition on, and which decisions should be delegated to the innovators who hold superior local information.

Our approach quickly narrows the range of choices from more than a dozen instruments to a small set where decision rights align with information distribution. We also characterize how economic factors such as financing constraints and coordination needs, along with practical implementation considerations, further refine the choice within this viable set. By working through this framework, a policymaker can efficiently identify a few well-suited mechanisms to evaluate and tailor to the specific innovation goal. 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 targeted mechanism recommendations for their specific innovation challenge.



FIGURE 1. SCREEN CAPTURE OF COMPANION WEB TOOL, *ATLAS OF INNOVATION*

The remainder of this paper proceeds as follows. Section 2 reviews the relevant literature on innovation policy, information economics, and mechanism design. Section 3 formalizes the innovation funding problem, characterizing how information distribution shapes contracting viability and efficiency. Section 4 presents our framework in detail, walking through the three core questions noted above, and mapping them to specific funding mechanisms. Section 5 concludes by discussing implications for policy and directions for future research.

2. Literature Review

The case for public innovation funding rests on well-established information frictions. 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 (Glennerster and Kremer 2004; Jaffe et al. 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 et al. 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 et al. 2012, Kremer et al. 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.

Multiple practitioner playbooks have cataloged the breadth of these funding mechanisms (Nesta n.d., Levers for Progress n.d., Renaissance n.d.), highlighting the benefits, drawbacks, and implementation details of each approach, and academic literature has

provided rigorous theoretical comparisons and toolkits (Wright 1983; Bloom et al. 2019; Azoulay and Li 2020; Armitage et al. 2023). In other policy domains, researchers have organized policy instrument choice around directly observable features of the policy environment rather than latent economic constructs (Hausmann et al. 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 a core insight from incomplete contracting theory: mechanism feasibility is strictly constrained by the distribution of verifiable information between principal and agent (Harris and Raviv 1979; Spier 1992; Lerner and Malmendier 2010).

3. Conceptual Model

3.1. Defining problems, solutions, and performers

The framework we develop centers on three dimensions of decision-making in public R&D programs: which problem to target, what constitutes a successful solution, and who will carry out the work. As we will show, these three dimensions jointly determine which mechanisms function well. We begin by defining the core concepts that shape our analysis.

An *innovation goal* represents the broad objective that motivates policy intervention. It is a desired transformation of the state of the world, from the current technological capabilities, scientific understanding, and social outcomes to an envisioned goal state, such as a world with new knowledge or new technology. Examples of such goals include 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), and the large-scale deployment of a product or service (such as a low-cost vaccine manufactured and distributed at scale). Innovation policy is concerned with increasing the probability that the world transitions from the current state to the goal

state. A realized innovation goal is composed of a problem and a solution pair, and a performer or many performers who identified successful technological approaches.

Within this framing, a *problem* identifies the gap, challenge, or specific obstacle that must be overcome to achieve the innovation goal. Problem statements vary in precision: broad articulations like "improve pandemic preparedness" leave many dimensions unspecified, while narrow articulations like "develop rapid vaccine manufacturing capacity for emerging pathogens" provide clear targets. Sufficiently narrow problem articulation is a prerequisite for defining solutions, as contractible success criteria cannot be specified without clarity about what problem is being solved.

A *solution* is a technological approach that successfully addresses the problem and achieves the innovation goal. While a solution cannot exist without a problem, the reverse is not true: a policymaker may be able to articulate a problem to fund without yet knowing what a solution would look like. Solutions can be characterized in terms of performance specifications and characteristics that an approach must satisfy. Output specifications establish verifiable, contractible criteria, such as efficacy thresholds, cost targets, or technical requirements, without dictating the technological approach for achieving them. For instance, in addressing the challenge of developing rapid vaccine manufacturing capacity, a policymaker might specify that a vaccine production platform must be capable of delivering 100 million doses within six months of pathogen identification and achieving at least 70% efficacy, leaving open which technological approaches developers pursue, whether mRNA, viral vector, or protein-based platforms. Specifying a solution requires that the problem be articulated narrowly enough that the success criteria are identifiable and measurable.

A *performer* is an individual or organization that selects and implements a specific scientific or technological approach in pursuit of the innovation goal. Performers possess specialized knowledge, technical expertise, and resources, and they hold private information about their own capabilities, the costs of different approaches, and beliefs over which approaches are likely to succeed.

A defining feature of innovation is that the set of approaches that would achieve the innovation goal is unknown to both policymakers and performers, creating bilateral uncertainty absent in standard procurement. In standard procurement, the policymaker knows what output they want and contracts with a performer to deliver it under a specified cost structure. If the performer fails to deliver, that failure reflects contractor performance; the project was achievable, but this particular performer did not execute. In innovation funding, failure has multiple potential sources; the performer may lack competence, the chosen solution specifications may not result in the goal, or the underlying problem may be fundamentally unsolvable given current knowledge and technology. This diffusion of risk across multiple dimensions — performer capability, solution feasibility, and problem tractability — makes innovation a more complex contracting challenge than standard procurement.

This conceptual framing motivates our focus on how decision rights over problems, solutions, and performers are allocated under informational constraints. We now formalize the policymaker's objective, before turning to the performer's objective.

3.2. The policymaker's objective

We take the perspective of a policymaker who has determined that public investment is needed to spur R&D, has been endowed with an innovation goal, and has a set budget to achieve it. We assume intervention is justified by the presence of some market failure — such as externalities, public goods characteristics, or coordination problems — that causes commercial markets to underinvest. We take this determination of a market failure as given, along with the policymaker's limited budget. As such, we set aside the wider landscape of innovation policies that are not directly connected to disbursing funds, such as high-skilled immigration, patent design, and regulatory reform (Bloom et al. 2019).

The policymaker's objective is to maximize the probability of achieving an innovation goal subject to a budget constraint and the available information. The probability of achieving an innovation goal depends on both how well the policymaker can describe their innovation goal — whether the right problems are targeted and whether solution criteria

are correctly defined — and how many capable performers the policymaker incentivizes to work on achieving the goal. The main way the policymaker incentivizes effort is through monetary payment or in-kind transfers to performers from their pre-specified budget.

To maximize the probability of achieving their innovation goal, the policymaker specifies the conditions of payment, whether immediate contracts or promises of future payment. The design features of these monetary agreements, or funding mechanisms, help determine which performers choose to engage and how much effort they invest.

The probability of achieving the innovation goal depends directly on how many performers attempt the innovation goal and each performer's own probability of success. In the cases where multiple performers attempt new scientific or engineering approaches, the aggregate probability of success for the policymaker depends on the correlation of the technical approaches within the portfolio. Some technical approaches to solving a problem share common failure modes, and when this is the case, funding more performer attempts does not diversify risk as much as when performers attempt uncorrelated, independent approaches.

A policymaker only has a chance at succeeding in their innovation goal if their incentives induce participation and effort from performers. We now turn to the performer's objective and how it is shaped by the policymaker's design choices.

3.3. The performer's objective

Performers are the agents who exert costly effort and make internal technical and organizational choices in pursuit of a solution. These choices include which scientific or engineering approaches to pursue, how to allocate resources, and how to respond to new information. Policymakers typically cannot observe these internal decisions directly and instead rely on intermediate outputs or outcomes that only imperfectly predict eventual success. Performers may hold private information about the feasibility of their approaches and their own capabilities, creating an information asymmetry that shapes which contractual arrangements are efficient.

Each potential performer decides whether to participate at all, and if so, how much effort to dedicate to an approach. A potential performer's objective is to maximize their net benefits, and performers decide to pursue an innovation goal when their expected net benefits exceed their next-best option. Benefits accrue from the monetary payment that the funding mechanism promises, along with any additional benefits of success, such as reputation or property rights. The performer incurs costs through the time, money, and effort spent attempting the innovation goal — and potentially additional costs upon failure. Depending on a policymaker's choice of mechanism, payment may or may not be conditioned on achieving features of the goal state.

Since success is inherently uncertain and potentially impossible, the performer weighs how likely a given approach is to succeed, how well they can execute that approach, and how expensive such an effort is. Performers differ in their capabilities, information about approach feasibility, and costs. This private information shapes a performer's entry decision and effort intensity. Conditional on participating, a performer chooses the approach and effort that best balances higher chances of success against higher costs, given the information available to them.

Performers choose which approaches to explore and their level of effort, while policymakers shape performers' payoffs by choosing the design features of a funding mechanism. The policymaker tailors the mechanism to their innovation goal, based on the information available to them at the outset. Both the policymaker and performer are constrained by the information available to each, respectively, at the time of their decision.

3.4. Distribution of information

Given the conceptual distinction between the aspects of an innovation goal (problems, solutions, and performers), the key determinant of mechanism viability is how information about each is distributed across actors. Information problems in innovation funding include both uncertainty and asymmetric information. Many aspects of program design are concerned with information revelation and create structures to draw out asymmetric information held only by the performer. However, absent asymmetry and under

uncertainty, there is no information yet to reveal, and the policymaker can contract on what knowledge they have despite the uncertainty. Three dimensions of information are core: Can the policymaker articulate which problem warrants investment? Can the policymaker define what success looks like in verifiable terms? Can the policymaker identify which performers are capable of executing?

Articulating the problem requires the policymaker to specify which problem domain is valuable with enough precision to condition funding on it. A policymaker may decide that climate change matters but still struggle to specify which problems within that domain should be prioritized. Researchers working at the frontier may possess superior knowledge of which technical problems, if solved, would be transformative. When the policymaker lacks this knowledge, delegating problem selection to performers becomes more attractive.

Determining what constitutes a solution requires knowing enough about the problem to specify success criteria. In some cases, policymakers can clearly define what is needed, such as a vaccine with a given efficacy threshold or a battery that meets specific performance benchmarks. In other cases, the policymaker may know a problem is important but lack the technical knowledge to specify what a solution would look like, or may not be able to verify that a proposed solution actually works, such as creating safe artificial intelligence systems. When verifiable success criteria cannot be specified well, or when the best available proxy diverges from the true goal, mechanisms that condition payment on measured outcomes will distort effort toward the criteria or proxy rather than the true goal.

The problem and solution dimensions both reflect the policymaker's information about the innovation challenge — and to some degree, they are jointly determined, since articulating a problem more precisely tends to clarify what a solution must look like. But the relationship is asymmetric: a policymaker may be able to articulate a specific problem without being able to specify a solution. These two dimensions capture different, although related, information: policymakers have the most information when they can articulate both the problem and the solution, and the least when they cannot articulate even the problem.

Evaluating which individuals, teams, or organizations can successfully execute a project requires assessing technical skills, organizational capacity, access to complementary expertise, and financial stability. Some of this information is observable through publication records, prior grants, and institutional affiliations. Performers typically hold superior information about their own capabilities, as well as their costs and their beliefs about which approaches are likely to succeed. They may also, however, have incentives to overstate capability. When a policymaker's signals about performer quality are weak, mechanisms that induce self-selection, where performers reveal private information through their willingness to participate under different contract terms, become more attractive.

The three dimensions vary in how uncertainty and asymmetric information arise across them. Performer capability is principally a problem of asymmetric information: performers hold private information about their own capabilities that the funder cannot directly verify. Problem articulation and solution specification can involve asymmetric information when frontier researchers hold superior knowledge unavailable to a generalist funder, but often also involve uncertainty inherent to innovation that neither side can resolve in advance.

Policymakers and innovators also each hold distinct pieces of information that are relevant for innovation success. Funding mechanisms differ in how they allocate decision-making authority across these actors. For example, project-based grants require specific answers to all three core questions, person-based grants specify the performer but leave problems and approaches open, and pull mechanisms define the problem and solution criteria but rely on competition to reveal capable performers.

Beyond these three core dimensions, policymakers may also have additional beliefs about features of the innovation environment that affect which mechanisms are attractive. These include: how long the path from investment to outcome is likely to take (time horizon); whether different technical approaches share common failure modes, limiting the value of funding parallel attempts (technical correlation); whether performers working on related problems would benefit from sharing information or resources (coordination effects); and how much performers differ from one another in capability (performer

heterogeneity). These are features of the world, and the policymaker's beliefs about them can shape mechanism choice. Misjudging these features can lead to mechanisms that exclude capable performers, fund superficially diversified portfolios, or miss gains from coordination. We discuss these features in section 4.2.

3.5. Funding mechanisms

Funding mechanisms are the link between the policymaker's objectives and performer incentives. A mechanism specifies what is paid, when payment occurs, and what conditions must be met. These design choices determine which performers participate, how much effort they invest, and which approaches they pursue.

TABLE 1—VARIETY OF PUBLIC FUNDING MECHANISMS

Mechanism	Definition
Advance market commitment (AMC)	Guarantee to subsidize or purchase at a specified price a yet-to-be-invented product meeting specified technical requirements.
Coordinated research programs	Institutional funding framework designed to accelerate high-risk, high-reward research, overcoming complex innovation challenges through active program management.
Fieldbuilding	Effort to create or strengthen entire disciplines, sectors, or professional ecosystems by training new talent, creating resources, and building legitimacy. Covers strategies such as hosting networking events, generating media interest, awarding recognition prizes, and supporting scholarship and apprenticeship programs.
Research joint ventures	Collaboration between a government agency and an external research institution, often formalized through agreements like Cooperative Research and Development Agreements (CRADAs). Share resources, knowledge, and expertise to tackle shared scientific challenges or develop breakthrough technologies.
Intramural science	Research conducted within government or philanthropic institutions, using in-house teams, often in government-owned facilities. Frequently involves sensitive or classified data and large-scale infrastructure.
Loan instruments	Action by public funders to reduce innovators' borrowing costs, either by directly providing loans on favorable terms, subsidizing interest rates on private loans, offering flexible repayment terms, or guaranteeing repayment to lenders if borrowers default.
Milestone payments	Staged funding disbursements rewarding incremental progress toward an innovation target. Innovators receive pre-specified payment tranches upon successfully achieving key technical or development milestones.
Per-unit subsidy or per-unit tax credit	Predetermined payment for each unit of a specified product or service produced or purchased. Creates volume-based incentives that scale with output or adoption.
Prize, absolute	Lump-sum reward for innovation meeting a technical target.
Prize, relative	Lump-sum reward for innovation outperforming peers'.
Procurement contracting	Agreement to purchase innovative goods or services meeting a specific technical or functional need.

R&D contracts	Agreement between funder and performer (often a company) to conduct R&D on a specified problem delivering a clearly defined idea, insight, proof of concept, or early-stage prototype within an agreed-upon timeline and budget.
R&D tax credit	Allows companies to reduce their tax bill by claiming credits on eligible R&D expenses.
Research grant, project-based	Funding given to individuals or organizations to cover research costs to support research on a specific research question with a defined method or goal. Often provided upfront or as cost reimbursements.
Research grant, researcher-based	Funding given to individuals or organizations to cover research costs for a broad research agenda. Often provided upfront or as cost reimbursements.

Notes: List in alphabetical order of mechanisms public funders commonly use to support innovation. Each can take a variety of forms; but for brevity we focus on typical forms of these agreements.

Table 1 defines the 15 funding mechanisms we examine throughout this paper. The definitions in the table reflect typical distinctions, but the boundaries between mechanisms are not always clear-cut, with substantial variation in implementation. These categories establish a vocabulary for our analysis, but do not constitute an exhaustive or mutually exclusive typology of innovation funding mechanisms.These funding mechanisms receive a more complete treatment in the companion *Atlas of Innovation* online tool (Institute for Progress and Market Shaping Accelerator 2026).

What a policymaker can viably write into a funding mechanism (which problems to target, which solutions to reward, which performers to select) depends on which of these dimensions the policymaker has information about. A policymaker can only contract well on dimensions they are knowledgeable about, either through observation or verification. If the policymaker cannot assess the value of different problems, the funding mechanism cannot condition payment on solving valuable ones and must instead delegate problem selection. If the policymaker cannot evaluate the feasibility of proposed solutions or specify verifiable success criteria, the funding mechanism cannot specify deliverables in advance and must allow for exploration and adaptation. If the policymaker cannot reliably identify performer capability, the funding mechanism cannot select performers directly and must instead design mechanisms that induce self-selection by capable parties.

Depending on a policymaker's choice of mechanism, payment is conditioned on or designed around descriptions and features of the goal state. The funding mechanism connects the policymaker's objective to the performer's objective by specifying the size and conditions of payment.

The spread between payoffs under success and failure drives performer behavior. Under some mechanisms, performers receive the same payment regardless of whether they reach the innovation goal; this small spread means the policymaker bears the financial risk of failure but provides weak incentives for effort. Under others, payment is contingent on success, concentrating risk on the performer but sharpening the link between effort and reward. Ideally, the policymaker picks a mechanism whose spread best aligns performer

incentives with the innovation goal, given what the policymaker knows and what performers can bear.

This logic generates the framework that follows. By characterizing the policymaker's ability to specify problems, solutions, and performers, we can identify which mechanisms are well-suited to different informational environments.

4. A Framework for Efficient Mechanism Selection

The NASA COTS example showed how different mechanisms produce different outcomes when applied to a single innovation goal — that mechanism choice matters for effective investment. Before presenting the framework, we turn to an example where a single mechanism produces different outcomes when applied to two technologies whose underlying characteristics differ. The information environment matters for selecting efficient funding mechanisms.

The Energy Independence and Security Act of 2007 directed U.S. fuel suppliers to blend increasing volumes of biofuels into their fuel. The policy applied the same instrument to two technologies, with divergent outcomes. For one of these technologies, corn ethanol, production scaled steadily and met its statutory volumes through 2022. For the other, cellulosic ethanol, production has lagged severely behind its statutory schedule, and EPA has been forced to revise down the schedule repeatedly. The divergence reflects the technological state of each industry at the time the mandate was written (Clancy and Moschini 2017). Corn ethanol relied on a mature production process where the outstanding problem was adoption. Guaranteed demand was sufficient to pull capacity into the market. Cellulosic ethanol still required scientific advances to make production scalable and commercially viable.

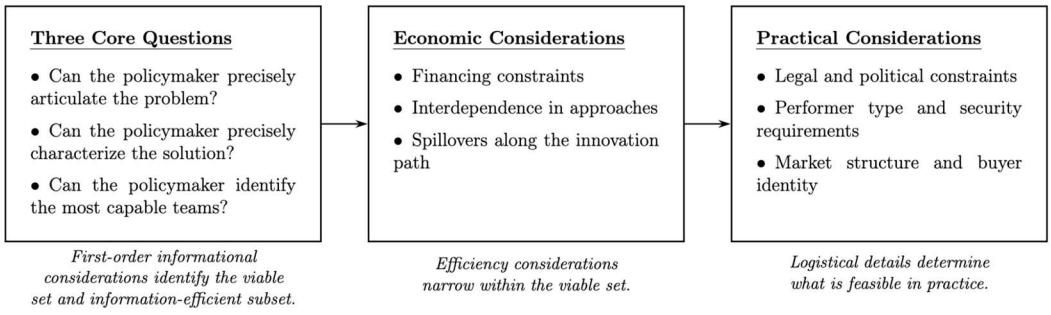


FIGURE 2. THREE-PART FRAMEWORK FOR FUNDING MECHANISM SELECTION

The framework that follows is built to make distinctions of this kind tractable to policymakers in advance of funding mechanism selection. This section develops the framework for selecting among funding mechanisms. The three core questions ask what the policymaker knows about the problem, solution, and performers, and identify the efficient and viable mechanisms. Economic considerations refine the viable set further with additional information available to the policymaker. Last, practical considerations determine what is feasible to implement. Figure 2 illustrates the framework, previewing the structure of this section: the three core questions narrow a set of 15 mechanisms to a viable subset, which is then refined by economic considerations (Section 4.2) and practical constraints (Section 4.3).

4.1. The three core questions

Our framework starts with three core questions. Each question operationalizes one dimension of the information structure about the innovation goal from section 3.1. Each corresponds to contractual viability and efficiency.

We use ‘efficient’ to mean information-aligned. A funding mechanism is efficient if it conditions funding, selection, or rewards on dimensions where the policymaker has

sufficiently accurate and verifiable information, and delegates decisions on other dimensions to innovators who hold superior private information.

Mechanisms that do not condition on dimensions where the policymaker does hold information are viable in principle, but inefficient in practice. Such mechanisms predictably reduce the likelihood of research success through adverse selection, moral hazard, or misallocation of effort. Hence, our definition reflects a second-best efficiency criterion: mechanisms are efficient if they align decision rights with information in a way that maximizes the probability of success given a policymaker's informational constraints.

These three questions are not complete in detailing 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 about the innovation goal they hope to solve. The three parts of the innovation goal — problem, solution, performers — interact with additional dimensions of the innovation environment, which we discuss in sections 4.2 and 4.3. These other features of the world that can constrain performers refine mechanism choice and design, but are not universally knowable or required for every innovation goal.

From the perspective of a policymaker, these three questions are part of defining (or not defining) aspects of an innovation goal and, as such, are likely to be answerable, perhaps more so than other economic elements, such as estimating the extent to which adverse selection is a problem.

We discuss each question in turn.

Can the policymaker precisely articulate the problem?.— The first question is whether a policymaker understands which problems are relevant to their innovation goal. Policymakers can have general objectives such as "advance the field of quantum computing" or "improve understanding of aging." These may be important goals, but they 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 delegate problem selection to performers with critical information advantages. When capable performers exist, they have greater knowledge of which specific questions are tractable, which

approaches are promising, and when serendipitous discoveries should change priorities. It is more effective to draw on research performers' expertise to determine which research paths to pursue by delegating problem selection. This precludes policymakers from writing contracts that are tightly specified in terms of deliverables or outcome-based payments in advance. Funding mechanisms available to policymakers in this situation include R&D tax credits, fieldbuilding, person-based grants, and intramural science. These more open and flexible instruments allow performers to set research directions and priorities. When the policymaker can neither articulate the problem nor identify performers, only R&D tax credits and fieldbuilding remain.

When policymakers can articulate problems precisely, they can contract on which questions are in scope, and allow for contracting on deliverables and results if they can specify outputs and verify when they have been achieved. This can enable the use of project-based grants and contracting, and opens the door to prizes, AMCs, and procurement.

Can the policymaker precisely characterize the solution?.— The second question is whether the policymaker knows enough about what a successful solution entails to specify it in a contract. This is distinct from understanding the problem. A policymaker may know precisely what problem 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 policymakers cannot characterize the solution, they cannot write performance-based contracts or define success criteria in advance. When a policymaker can articulate the problem but not the solution, technical uncertainty is high, 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 instruments allow researchers to pivot as new information emerges and to pursue

approaches the policymaker could not have anticipated. They fund the effort on a problem rather than the success of a performer's chosen approach.

When a policymaker can characterize the features a successful solution must possess, contracts can condition on outcomes by including target product profiles, performance specifications, and minimum efficacy thresholds. Success becomes observable and verifiable even if no solution yet exists. The technical path remains flexible, since innovators choose their approach, but the endpoint is fixed. Specifying the solution criteria also allows policymakers to be firm- and technology-agnostic. Mechanisms that pay for results include absolute prizes, AMCs, milestone payments, and per-unit subsidies. These mechanisms rely on specifying success criteria up front and conditioning payment on achieving them.

Can the policymaker identify the most capable teams?.— The third question is whether the policymaker has sufficient information to select R&D performers directly or must delegate selection to the mechanism itself. This combines two related information problems: which technical approaches are most promising, and who is best positioned to deliver them. Information on R&D performers can be acquired through auctions or requests for proposals, but 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 three conditions hold. First, a performer pool must exist. Without trained researchers or capable firms, there are no performers to select among, and fieldbuilding must precede project funding. Second, performers must differ meaningfully in capability. If all potential performers are roughly equivalent, costly selection processes do not improve outcomes. Third, the policymaker must be able to distinguish capable from incapable performers using observable signals that correlate with capability, such as track records, proposals, or peer review. When all three conditions are met, grants and contracts that rely on direct selection become more attractive.

When policymakers cannot reliably distinguish the most capable teams, innovators hold private information about their own capabilities and costs that they cannot credibly communicate to policymakers. 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 the 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 projects.

TABLE 2—MECHANISM VIABILITY BY CORE QUESTION

Answer to core question			
Articulate problem?	**Characterize solution?**	**Identify teams?**	**Viable mechanisms (efficient starred)**
No	Not applicable	No	Fieldbuilding* R&D tax credits*
No	Not applicable	Yes	Fieldbuilding, R&D tax credits, Researcher-based research grants* Intramural science*†
Yes	No	No	Fieldbuilding R&D tax credits Relative prizes*
Yes	No	Yes	Fieldbuilding R&D tax credits Researcher-based research grants Intramural science Coordinated research program* Government–research lab joint ventures* R&D contracts*
Yes	Yes	No	Fieldbuilding R&D tax credits Relative prizes Absolute prizes* Milestone payments* Advance market commitment* Per-unit subsidies*
Yes	Yes	Yes	Fieldbuilding R&D tax credits Researcher-based research grants Intramural science Coordinated research program Government–research lab joint ventures R&D contracts Relative prizes Absolute prizes Milestone payments

	Advance market commitment
	Per-unit subsidies
	Procurement contracting*
	Project-based research grants*
	Loan subsidies*

Notes: Table lists viable funding mechanisms for each possible configuration of answers to core questions. Rows lower in the table require more information, so "yes" answers unlock more viable mechanisms. Those mechanisms that exploit the expanded information available to the funder in lower rows are considered "efficient" and designated with stars.

† Intramural science fits here in its open-ended form, where in-house researchers set their own agendas within a broad mission area. Some intramural programs are better motivated by security requirements or long-term institutional needs than by uncertainty about the problem, which are considerations that appear elsewhere in the framework.

Table 2 classifies mechanisms by the three core questions, treating each as a yes/no judgment. In practice these are matters of degree, but the binary framing here makes the table tractable. Mechanisms marked with an asterisk are efficient: they use all available information optimally given the policymaker's answers. Reading down the table, the viable set expands: fieldbuilding and R&D tax credits remain available throughout, while mechanisms like procurement contracting and project-based grants only become viable when the policymaker can answer all three questions affirmatively. Additional considerations in sections 4.2 and 4.3 help narrow the viable set for reasons the three core questions do not capture.

4.2 Economic Considerations

The three core questions operationalize familiar concepts from information economics and innovation policy. Whether a policymaker can articulate the problem and specify the solution reflects the technological maturity of the domain. Performer identification addresses 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. Other considerations from innovation economics, such as appropriability, coordination needs, and financing constraints, are relevant but operate within the viable set identified by these three questions (the second narrowing step in Figure 2).

Within the viable set, choosing well requires understanding how mechanisms will function in practice. These factors fall into two categories: economic considerations that affect mechanism performance, and practical constraints that limit implementation.

The economic considerations discussed in this section are secondary refinements that operate within the viable set of mechanisms identified by the three core questions. They can matter significantly for mechanism design and fit.

Financing constraints.— Mechanisms that defer payment, such as pull mechanisms that condition payment on successful delivery or push mechanisms structured as reimbursements, require performers to finance their own effort until outputs are verified. This financing burden grows with the time horizon between investment and payment. When projects require years of development before outcomes can be assessed, 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. 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 (such as startups, university labs, and firms in developing economies) 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 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.

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 appears. 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 advance market commitments. When the innovation task is relatively straightforward and duplication 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.

Spillovers along the innovation path.— Innovation often occurs in stages — initial invention followed by scale-up and commercialization — potentially performed by different actors. If early-stage innovations are appropriable through intellectual property, later-stage performers can compensate earlier innovators through licensing, allowing mechanisms to focus solely on outcomes. When early-stage contributions generate non-appropriable spillovers, lessons spread to the whole field rather than staying with the originator, and mechanisms tied to outcomes will inadequately incentivize the crucial first steps. In these cases, mechanisms must explicitly reward each stage. Milestone payments can address this by providing intermediate rewards for breakthroughs that enable subsequent innovation, even when those breakthroughs cannot be protected or sold.

4.3. Practical Considerations

Not every mechanism that is viable in principle is feasible for policymakers in practice. The next several considerations may force policymakers away from economically optimal choices (the third narrowing step in Figure 2).

For example, ARPA-style programs require empowered program managers with technical expertise and authority to make rapid funding decisions, while prizes require the capacity to specify and verify achievement criteria. Advance market commitments and prizes require institutional infrastructure to credibly commit to future payments. Agencies lacking these capabilities must choose mechanisms they can actually implement, even if other mechanisms would be theoretically preferable.

Legal and political constraints.— Legal authority constrains what agencies can do. Statutory mandates and procurement regulations may rule out mechanisms that would otherwise be attractive. Political considerations matter as well: budget scoring conventions, time constraints, or oversight risk may shape mechanism choice. For example, mechanisms that concentrate visible resources on a small number of performers may be harder to justify and sustain politically than mechanisms that distribute funding broadly, even if concentration would be more effective.

Performer type and security requirements.— Beyond identifying whether capable performers exist, the specific type of performer best suited to the innovation goal further narrows mechanism choice. Government intramural labs are well-suited for problems requiring government information, especially classified data, and classified work may require intramural execution or tightly controlled contractor relationships. Government labs may also be appropriate for long-term programs requiring institutional stability or when regulatory neutrality is a relevant concern.

Academic researchers are well-suited to create knowledge, particularly through long-term fundamental research, and for its dissemination and training. Academic researchers are less well-suited to commercializing their research and likely less responsive to commercial rewards.

Private-sector firms have strengths in bringing products to market, scaling production, and responding to market incentives. Leveraging these strengths requires firms to be able to capture some of the value their innovations generate.

The choice of performer type constrains mechanism selection: mechanisms must be compatible with the capabilities and incentive structures of the target performer pool.

Market structure and buyer identity.— When the solution is a product, the identity of the primary buyer affects which mechanisms are feasible. If the government is the main buyer, procurement contracts become the natural choice, as the government can specify requirements and guarantee purchase. If a commercial market exists, loan instruments and per-unit output subsidies likely dominate, as private buyers provide revenue streams that can support repayment or justify subsidized risk-taking. Advance market commitments can serve as a middle ground, specifying a future per-unit subsidy and potentially including a guaranteed government purchase to establish a market floor. This approach works when the government wants to stimulate commercial production for goods with broader social benefits.

In an innovation funding environment with multiple public funders, international actors, or private investment, mechanism choice should account for what others are already funding: a prize may be redundant if other agencies already provide grants for the same goal, or it may complement existing support by rewarding outcomes where others fund exploration. International competition may make certain mechanisms less attractive if other jurisdictions offer better terms. Private capital markets may fill gaps, reducing the need for certain public mechanisms, or may amplify specific approaches that public funding could counterbalance.

5. Conclusion

We propose an information-based framework for selecting among innovation funding mechanisms. Rather than ranking mechanisms by expected value (a task that is rarely tractable given the uncertainties involved), we organize mechanism choice around three questions about the policymaker's information environment: Can the policymaker precisely articulate the problem domain? Can the policymaker precisely characterize the solution? Can the policymaker identify the most capable performers? These questions

determine which mechanism classes are feasible and which decisions must be delegated to better-informed innovators.

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 would raise success probability, and argues against naive diversification when approaches share common failure modes. When there are non-appropriable spillovers along the innovation path, the mechanisms must sufficiently compensate for them to ensure investment in each innovation stage.

The framework provides practical guidance grounded in economic principles under some simplifying assumptions. The framework assumes a market failure and that a policymaker has appropriately weighed the costs and benefits of intervention. As such, the framework can inform how to pursue a goal once chosen, but not which goals deserve public funds in the first place, nor the amount of public funds that should be committed to an innovation goal.

Our framework focuses on *ex ante* considerations. *Ex post* evaluation of program performance complements the *ex ante* reasoning developed here. Program implementation reveals information about program design that was unavailable *ex ante*, and *ex post* evaluations of similar past programs can complement the *ex ante* reasoning offered here. Empirical evaluation and randomized experimentation built into early program pilots allows policymakers to draw on both *ex ante* reasoning and *ex post* evidence during program implementation.

Some market failures may be better corrected by assigning property rights or changing statutes or regulations, rather than by deploying funding. The framework does not resolve all policymaker choices regarding innovation funding, but it offers a structured way for a policymaker to think through how to effectively deploy funds toward an innovation goal. Decisions about intellectual property protections, immigration policy for high-skilled workers, and regulatory reform represent a distinct class of innovation levers that operate

outside the funding mechanism problem that we address in this paper. Whether and how policymakers should choose between deploying funds versus assigning property rights is an understudied question, and future research on that allocation decision would be a valuable contribution.

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.

REFERENCES

Akerlof, George A. 1970. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism.” *The Quarterly Journal of Economics* 84 (3): 488–500. https://doi.org/10.2307/1879431.

Armitage, Sarah, Noël Bakhtian, and Adam B. Jaffe. 2023. “Innovation Market Failures and the Design of New Climate Policy Instruments.” In *Environmental and Energy Policy and the Economy, Volume 5*. University of Chicago Press. https://www.nber.org/books-and-chapters/environmental-and-energy-policy-and-economy-volume-5/innovation-market-failures-and-design-new-climate-policy-instruments.

Arrow, Kenneth. 1962. “Economic Welfare and the Allocation of Resources for Invention.” In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton University Press. https://www.nber.org/books-and-chapters/rate-and-direction-inventive-activity-economic-and-social-factors/economic-welfare-and-allocation-resources-invention.

Arrow, Kenneth. 1963. “Uncertainty and the Welfare Economics of Medical Care.” *The American Economic Review* LIII (5).

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). https://doi.org/10.1086/699933.

Azoulay, Pierre, and Danielle Li. 2020. “Scientific Grant Funding.” Working Paper No. 26889. Working Paper Series. National Bureau of Economic Research, March. https://doi.org/10.3386/w26889.

Azoulay, Pierre, Joshua S. Graff Zivin, and Gustavo Manso. 2011. “Incentives and Creativity: Evidence from the Academic Life Sciences.” *The RAND Journal of Economics* 42 (3): 527–54.

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. https://doi.org/10.1111/joes.12074.

Bloom, Nicholas, John Van Reenen, and Heidi Williams. 2019. “A Toolkit of Policies to Promote Innovation.” *Journal of Economic Perspectives* 33 (3): 163–84. https://doi.org/10.1257/jep.33.3.163.

Bloom, Nick, 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. https://doi.org/10.1016/S0047-2727(01)00086-X.

Brunt, Liam, Josh Lerner, and Tom Nicholas. 2012. “Inducement Prizes and Innovation.” *The Journal of Industrial Economics* 60 (4): 657–96. https://doi.org/10.1111/joie.12002.

Clancy, Matthew, and GianCarlo Moschini. 2017. "Mandates and the Incentive for Environmental Innovation." *American Journal of Agricultural Economics*. https://doi.org/10.1093/ajae/aax051.

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. https://doi.org/10.1257/pol.20200739.

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. https://doi.org/10.1016/j.respol.2018.05.005.

Harris, Milton, and Artur Raviv. 1979. “Optimal Incentive Contracts with Imperfect Information.” *Journal of Economic Theory* 20 (2): 231–59. https://doi.org/10.1016/0022-0531(79)90073-5.

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. https://doi.org/10.1093/acprof:oso/9780199534081.003.0015.

Holmström, Bengt. 1979. "Moral Hazard and Observability." *The Bell Journal of Economics* 10 (1): 74–91. https://doi.org/10.2307/3003320.

Howell, Sabrina T. 2017. "Financing Innovation: Evidence from R&D Grants." *American Economic Review* 107 (4): 1136–64. https://doi.org/10.1257/aer.20150808.

Institute for Progress and Market Shaping Accelerator. 2026. *Atlas of Innovation*. 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*, Technological Change and the Environment, vol. 54 (2): 164–74. https://doi.org/10.1016/j.ecolecon.2004.12.027.

Kremer, Michael, and Rachel Glennerster. 2004. *Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases*. Princeton University Press. https://doi.org/10.2307/j.ctt1dr365r.

Kremer, Michael, Jonathan Levin, and Christopher M. Snyder. 2020. "Advance Market Commitments: Insights from Theory and Experience." *AEA Papers and Proceedings* 110 (May): 269–73. https://doi.org/10.1257/pandp.20201017.

Lerner, Josh, and Ulrike Malmendier. 2010. "Contractibility and the Design of Research Agreements." *American Economic Review* 100 (1): 214–46. https://doi.org/10.1257/aer.100.1.214.

"Levers for Progress." https://leversforprogress.com.

Li, Danielle. 2017. "Expertise versus Bias in Evaluation: Evidence from the NIH." *American Economic Journal: Applied Economics* 9 (2): 60–92. https://doi.org/10.1257/app.20150421.

Nordhaus, William D. 2004. "Schumpeterian Profits in the American Economy: Theory and Measurement." Working Paper No. 10433. Working Paper Series. National Bureau of Economic Research, April. https://doi.org/10.3386/w10433.

Pauly, Mark V. 1968. "The Economics of Moral Hazard: Comment." *The American Economic Review* 58 (3): 531–37.

Renaissance Philanthropy. "Playbooks." https://www.renaissancephilanthropy.org/playbooks.

Spier, Kathryn E. 1992. "Incomplete Contracts and Signalling." *The RAND Journal of Economics* 23 (3): 432–43. https://doi.org/10.2307/2555872.

Technology Readiness Assessment Guide: Best Practices for Evaluating the Readiness of Technology for Use in Acquisition Programs and Projects. 2020. GAO-20-48G. U.S. Government Accountability Office. https://www.gao.gov/assets/gao-20-48g.pdf.

United States Government Accountability Office (GAO). 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. https://www.gao.gov/assets/gao-09-618.pdf.

Vicki Sellick, Amy Solder, and Isobel Roberts. 2018. *Funding Innovation: A Practice Guide*. Nesta. https://media.nesta.org.uk/documents/Funding-Innovation-Nov-18.pdf.

Wright, Brian D. 1983. "The Economics of Invention Incentives: Patents, Prizes, and Research Contracts." *The American Economic Review* 73 (4): 691–707.