

BUILDING A GOOD JOBS ECONOMY*

Dani Rodrik

Charles Sabel

Harvard Kennedy School

Columbia Law School

Draft, April 2019

For comments only

* We thank Tim Bartik, Josh Cohen, and William Simon for very useful comments on an earlier draft.

1. Introduction and statement of the problem

Conventional models are failing throughout the world. In the developed world, the welfare state-compensation model has been in retrenchment for some time, and the drawbacks of the neoliberal conception that has superseded it are increasingly evident. Yet there is no compelling alternative on offer. In the developing world, the conventional, tried-and-tested model of industrialization has run out of steam. In both sets of societies a combination of technological and economic forces (in particular, globalization) is creating or exacerbating productive/technological dualism, with a segment of advanced production in metropolitan areas that thrives on the uncertainty generated by the knowledge economy co-existing with a mass of relatively less productive activities and communities that neither contributes to nor benefits from innovation. The sizes of these two sectors and the trajectories leading into them may vary, but otherwise the nature of the underlying problem seems to have converged in the developed and developing worlds.

This productive/technological dualism is in turn responsible for many of the ills these societies face: inequality, exclusion, spatial/social segmentation, loss of trust in elites/governments/experts, the populist backlash, and authoritarian politics. Left to their own, globalization and new technologies look likely to aggravate these divisions and the pathologies that flow from them.

Much of our policy conversation today focuses on solutions that elide the true source of the problem. Ex post redistribution through taxes and transfers accepts the productive structure as given, and merely ameliorates the results through handouts. Investments in education, universal basic income (UBI), and social wealth funds seek to enhance the endowments of the workforce, without ensuring productive integration. Broadly speaking the same can be said about the Keynesian approach to job creation, through aggregate demand management. Keynesianism aims at static efficiency—closing the gap between actual and potential output where the potential output is fixed enough to be precisely calculable. Dualism entrenched enough to shape long-term growth expectations—just the kind of structural deformation of the economy that most concerns us—cannot be addressed by demand management, short, perhaps, of mobilization for war. Though lax enforcement of anti-trust laws may have contributed to the concentration of industry in recent decades, and exacerbated inequality by allowing oligopolists to increase their markups or use monopsony power in labor markets to drive down wages, redress through a new round of trust busting is at best a very partial solution to the larger problem, and then only in the long term.

What we seek to explore here is a set of interventions by the public sector – or its delegated agencies – directly in the productive sphere, and in direct collaboration with the most productive segments of the private sector. These interventions are targeted at expanding productive employment opportunities by supporting firms and workers in their efforts to acquire and extend the capacities needed to participate in the dynamic sector of the economy. We call it a strategy of “building a good jobs economy.” It has three, mutually re-enforcing components: increasing the skill level and productivity of existing jobs, and the competitiveness of firms, for example through provision of extension services to improve management or cooperative programs to advance technology; increasing the number of good jobs by

supporting startups, the expansion of existing, local firms or attracting investment by outsiders—what the many state and local programs (of greatly varying quality) currently directed to this last purpose refer to simply as “economic development”; and active labor market policies or workforce development programs to help workers, especially from at-risk groups, master the skills required for good jobs. Redistribution, Keynesian demand management and anti-trust policies can and should be important complements to such interventions; alone or in combination they cannot be a substitute for them.

Public-private collaborations are at the heart of this strategy. Our focus is on the design principles needed to govern these collaborations. Such principles do not need to be invented from scratch. We argue that they can be borrowed from existing innovative governance arrangements firms, regulators and other public agencies have already developed in response to the market and technological uncertainties they face. These arrangements have not been typically deployed in pursuit of good jobs, but they can be adapted to that end.

Active labor market policies have begun to receive considerable attention. Spurred by that interest, a number of studies have recently carried out more encompassing reviews of development experience on the ground with some resemblance to the good jobs program. Austin et al. (2018) survey “place-based” policies such as regionally targeted employment subsidies and infrastructure investment, Miller-Adams et al. (2019) review program focused on creating good jobs in communities that face economic challenges, and Autor et al. (2019) summarize evidence on the impact of educational and other interventions. A common theme that emerges from this literature is the inter-related and conditional nature of the remedies; very few program elements work off the shelf and reliably across diverse settings. For example, Miller-Adams et al. (2019) recommend differentiated strategies that combine skill development strategies (targeted at the local workforce) with programs to attract businesses (targeted at employers). They emphasize the provision of customized services to firms and design elements that are tailored to local conditions.

In view of the inherent uncertainty about “what works,” we focus here not on specific interventions or policies, but on a meta-intervention regime that is structured to generate the requisite solutions. The objective is to provide a set of general principles for regimes that produce good jobs in many different areas of economic activity. We describe a set of design principles for building dynamic governance arrangements that sustain public-private collaborations under conditions of uncertainty and learning, through ongoing review and revision of objectives, instruments, and benchmarks.

The most familiar variant of a “meta-regime” is collective bargaining, or some close social partnership or corporatist analogue. The expectation is that, given a secure place in the constitutional order, public and private actors will bargain to achieve public-regarding outcomes. Responding both to the burdens of continuing adjustment imposed by uncertain and diverse conditions, as well as the chastening lessons of corporatist self-dealing, we found the meta-regime on governance principles subjecting all decisions to continuing mutual scrutiny by stakeholders under public oversight. These innovative modes of governance, we argue further, allow the parties, beginning with only a thin understanding of the substance and scope of their goals, to assess one another’s capacities and good faith in the very process of refining ideas of what the eventual project should be. Trust and coalition building—the preconditions

for meta-regimes built on social partnership—are the outcome of joint problem solving under this governance. The good-jobs strategy can only succeed in the end with the support of a wide and robust coalition. We argue that building consensus through problem solving helps ensure that mobilization is put to effective use.¹

We proceed as follows. First, we argue that the shortfall in “good jobs” can be viewed as a massive market failure – a kind of gross economic malfunction, and not just a source of inequality and economic exclusion. Next, we make the case that this problem cannot be dealt with standard regulatory instruments such as Pigovian subsidies. Pervasive uncertainty, dependence on differentiated local conditions, and the evolving nature of the goals call for a high-dimensional policy space and an iterative model of strategic collaboration between private actors and the state. We illustrate the organizational framework we have in mind using examples from two analogous policy domains where the principles are already in application: fostering of advanced technologies (DARPA and ARPA-E) and environmental regulation of (dairy farming in Ireland).

We then sketch out how these ideas can be applied in principle to the specific domain of good jobs, and briefly review their actual application to the three components of the strategy. The focus is on workforce development, where promising and arguably generalizable outcomes have been achieved under governance principles closely approximating those set out here. By way of conclusion we caution against the intellectual pessimism about job creation strategies that haunts many careful reviews of program evaluations, and suggest a reason for optimism with regard to the formation of the political will needed to build a good jobs economy. We should not make too much of disappointing reviews of program performance because, some recent studies aside, those reviews and the evaluations on which they are based do not discriminate between one-size-fits all programs and more successful programs that adapt policy to changing local circumstance. We see cause for some optimism on political will in the possibility that once in limited operation the governance arrangements we propose may help create and enlarge the constituency for acting on the problems they address.

2. *“Good jobs” as a source of positive externality*

The sources of the externality

The definition of “good job” is necessarily slippery. It depends on a host of contextual features: the overall level of productivity and economic development, costs of living, prevailing income gaps, and so on. Broadly speaking, we have in mind stable formal-sector employment that comes with core labor protections such as safe working conditions, collective bargaining rights, and regulations against arbitrary dismissal. A good job enables at least a middle-class existence, by a region’s standards, with

¹ We would argue that any broad program of mobilization, such as the green new deal, would have to rely on a similar mechanism to the extent that it involves uncertainty and contextualization.

enough income for housing, food, transportation, education, and other family expenses, as well as some saving.

We begin by establishing that producing good jobs is a source of positive externality for society. From an economic standpoint, the issues are analogous to those that arise in the cases of environmental externalities or R&D externalities, two domains on which we will draw when we develop our organizational recommendations.

A firm considers labor as a production input, with the market wage as its cost. In the short run, the wage rate determines the firm's desired level of employment. In the medium-run, it also determines the kind of technologies the firm invests in and the production technique – the mix between labor and various forms of capital. When wages rise, either because of greater productivity or enhanced bargaining power of labor, firms try to economize on the use of labor and adopt technologies that replace workers. From a society's standpoint, the result is an undesirable trade-off between good jobs and the level of employment. Today's economies tend to manage this trade-off by allowing dualistic labor markets to become entrenched (Temin 2017): islands of productive, high-wage activities exist in a sea of poor jobs. Labor market and social policies generally determine the distance between working conditions in the two sectors. But a higher floor on economy-wide wages generally comes at the expense of higher unemployment and lower labor hours.

Some version of this trade-off has existed throughout history. In growing economies the tension is typically alleviated by an economy-wide rise in productivity, which suppresses the distinction between insider and outsider jobs. For example, the mechanization of agriculture during the 19th and early 20th centuries created a surplus of labor in the countryside. But the workers who flooded into urban centers could be absorbed into manufacturing activities (and related services) where productivity and wages were even higher. De-industrialization during the second half of the 20th century led to a similar, but more challenging dilemma. Rapid labor productivity growth in manufacturing (and import competition) resulted in a loss of production jobs and a shift to employment in services, where wages and employment conditions were often inferior. Today's technological trends – automation, the knowledge economy, digital technologies – are leading to a significant exacerbation of the problem. The productivity effects of these new technologies remain bottled in a limited number of sectors and metropolitan locations, generating relatively small numbers of good jobs, while the rest of the economy remains stagnant (Remes et al., 2018). “Where will the good jobs come from?” is perhaps the defining question of our contemporary political economy.

We do not view this simply as a problem of inequality and exclusion, but also as a problem of gross economic inefficiency – a case of operating deep inside the production possibility frontier. That is because a shortage of good jobs is associated with a significant range of public ills. In economic terms, good jobs are a source of positive externalities.

The central distinction in an externality is between private and social costs. The private cost of labor is the market wage that an employer pays to a worker, net of employment taxes (or subsidies). The social cost of labor is its social opportunity cost, the value of output forgone in the rest of the economy when

the employer hires that worker. There are many things that could drive a wedge between the private and social costs. If the next best alternative for the worker is not to be employed at all, the social opportunity cost of labor could be very small (essentially just the personal disutility associated with work). If the alternative is a less productive and hence lower-paying job, the social opportunity cost will be higher, but still lower than the private cost.²

Such wedges due to involuntary unemployment or wage differentials, as important as they may be, are only a small part of the problem. The broader issue is the undermining of social structures that underpin economic prosperity. Communities where middle-class jobs have gone scarce suffer from a variety of social ailments.

In his path-breaking book When Work Disappears, sociologist William Julius Wilson (1996) described at length the social costs of the decline in the number manufacturing and blue-collar jobs, ranging from broken families to drug abuse and crime. While Wilson's focus was on racial minorities living in inner-city ghetto neighborhoods, his argument applies more broadly. Autor, Dorn, and Hanson (2018) studied communities across the entire U.S., differentiating them by the degree to which they were affected by import competition with China. Communities where jobs came under greatest pressure from Chinese imports experienced an increase in "idleness" among young males (the state of being neither employed nor in school) and a rise in male mortality due to drug and alcohol abuse, HIV/AIDS, and homicide. The adverse effects of job loss were not limited to men. There was also an increase in the fraction of mothers who are unwed, of children in single-headed households, and of children living in poverty.

These economic and social impacts of good jobs going scarce are compounded in turn by the political consequences. There is by now considerable evidence from a number of advanced market economies that links the rise of nativist populist political movements to adverse labor market developments. In the United States, Autor et al. (2017) have shown that the China trade shock had a significant impact on political polarization. Holding constant initial political conditions in 2002, districts that experienced sharper increases in import competition were less likely to elect a "moderate" legislator in 2010. New legislators elected in hardest hit areas tend to occupy more extreme positions on the ideological spectrum, especially on the right. Districts initially in Republican hands were substantially more likely to elect a GOP conservative. What is perhaps the most intriguing implication of this research is that the labor market disruptions due to the China trade shock may have been directly responsible for Donald Trump's electoral victory in 2016. Autor et al. undertake a counterfactual analysis in which they assume the growth of Chinese import penetration is 50 percent lower than the realized rate over the 2002-2014 period. Their estimates on the electoral consequences indicate that a Democrat instead of a Republican presidential candidate would have been elected in 2016 in the swing states of Michigan, Wisconsin, and Pennsylvania. The Democratic candidate would also have obtained an overall majority in the Electoral College under this counterfactual scenario.

² Austin et al. (2018) consider three sources of economic externalities from non-employment: fiscal costs on the state through the tax-transfer system, costs imposed on the family, and spillovers that encourage non-employment by others in the community. They reckon these costs range 0.21-0.36 times the wage of low-income workers. Our focus here is on social and political costs that we believe are much higher.

A recent paper on Sweden traces out very similar political consequences, even though the shocks that led to labor market disruption were of a different nature (Dal Bò et al., 2018). A series of reforms after 2006 under a Conservative-led coalition reduced social insurance and transfer benefits while lowering taxes, increasing the disposable income gap between “insiders” and “outsiders” – those with steady jobs and those who were either unemployed or relied on temporary jobs. The post-2008 financial crisis and recession further contributed to the widening of the gap. The main beneficiary appears to have been the right-wing, anti-immigrant Sweden Democrats party. The authors show that the local insider-outsider income gaps and the share of vulnerable insiders are positively correlated with larger electoral gains by the Sweden Democrats. Exposure to immigrants, on the other hand, is not systematically associated with support for the political right. The fundamental cause of nativist politics seems to be decline in secure, good jobs rather than cultural or xenophobic preferences per se.

Similar results have been reported for other European countries. Analyzing the political realignment behind Brexit, Colantone and Stanig (2016) attribute a key role to the labor market impact of globalization. Using an Autor et al.-type China trade shock variable, they show regions with larger import penetration from China had a higher Leave vote share. They corroborate this finding with individual-level data from the British Election Survey that shows individuals in regions more affected by the import shock were more likely to vote for Leave, conditional on education and other characteristics. A second paper by Colantone and Stanig (2017) undertakes a parallel analysis for fifteen European countries over the 1988-2007 period. It finds that the China trade shock played a statistically (and quantitatively) significant role across regions and at the individual level. A larger import shock was associated with support for nationalist parties and a shift towards radical right-wing parties. Guiso et al. (2017) look at European survey data on individual voting behavior and find an important role for economic insecurity – including exposure to competition from imports and immigrants – in driving populist parties’ growth. Individuals who experience greater economic insecurity were also less likely to show up at the polls.

Perhaps the most concerning aspect of the political consequences of adverse labor market shocks is that such shocks weaken support for democracy and foster authoritarian attitudes. The association between economic crisis and the rise of fascism in interwar Europe is well known (Frieden, 2006). More broadly, economic stagnation or decline among the middle classes undermines the set of moral values and beliefs that sustain liberal democracy (Friedman, 2005). There is evidence from our current moment in history that some of the same tendencies are at play. In the United States, individuals located in local labor markets that were more substantially affected by imports from China appear to have developed more authoritarian values (Ballard-Rosa et al., 2018). Similarly, individuals living in European regions that received more negative globalization shocks were systematically less supportive of democracy and liberal values and more in favor of authoritarian leaders (Colantone and Stanig, 2018).

In short, there are significant economic, social, and political costs of failure to generate good jobs. These costs drive a large wedge between the market wage and the social cost of labor. The social opportunity cost of generating good jobs is considerably lower than the market wage. Bad jobs lead to lagging communities with poor social outcomes (health, education, crime) and social and political strife (populist backlash, democratic malfunction). A private employer fails to take these costs into account, unless prompted to do by the state. The empirical literature suggests that these externalities are

substantial—perhaps so great that they threaten the economic order underpinning our form of government.

Our focus on the social externality of good jobs is a key difference from approaches that revolve around firm-level practices. For example, in her well-known book [The Good Jobs Strategy](#), Zeynep Ton advocates a range of employment policies such as higher wages and benefits that she argues could help employers as well as employees. The argument, nicely encapsulated in her subtitle, is that smart companies can boost profits by investing in their employees. We do not deny that such opportunities exist, and that firms may do well by doing good for their workers. But as Osterman (2018) emphasizes in a recent review, the evidence that profit maximizing firms can benefit from “high road” employment practices is limited and far from overwhelming. The vast majority of firms may not be inclined to offer or expand good jobs unless the strategy is part of a concerted collaboration with public agencies in which they are offered something in return. That something could be either carrots in the form of tailored public services or the withholding of a stick in the form of tax easements.

The inadequacy of standard remedies

Having established that good jobs are a source of positive externalities, we now discuss why standard remedies are inadequate.

The conventional instrument for internalizing an externality is a Pigovian subsidy, which would be a generalized employment subsidy in this case. But successful administration of Pigovian subsidies requires sufficient information about the size of the externality and (what often amounts to the same thing) a relatively static environment. In a dynamic environment with substantial uncertainty, alternative regulatory arrangements are often preferable.

In a classic article, Weitzman (1974) showed that quantity targets may dominate price instruments (such as a subsidy) under such conditions. A price instrument (subsidy) minimizes the costs of achieving a certain target, at the risk of missing the target (because of uncertainty about supply and demand responses, say). Quantitative targets, on the other hand, achieve the requisite social outcome, but potentially at greater economic cost than is necessary. When the risks of just missing the socially optimal target—making water drinkable, say—outweigh the risks of inadvertently imposing too high a clean-up cost on producers, quantity targets are preferable to Pigovian price instruments. The analogous argument in the present context is that the risk of failing to generate a sufficient number of good jobs in a particular community may dwarf the risk of imposing too high a burden on individual firms.

Uncertainty also increases the dimensionality of the policy space. In the standard conception of externalities, there is a single quantity, with an associated market price, that is responsible for the generation of the externality. The appropriate intervention consists of directly targeting that price (or quantity), and doing no more than that. But when there is uncertainty about behavior, technology, and the effectiveness of different policies, optimal policies – in the second-best sense of the term – will range over multiple margins of intervention and several different types of policy instruments. Learning

about what works and what does not becomes an integral part of the policy process. Establishing mechanisms of feedback from firms to public authorities is critical to the regulatory apparatus. The relevant policy space is of much higher dimensionality.

Finally, an additional problem with standard regulatory remedies in the present setting is that they postulate clear goals (“objective functions,” in economics jargon). As uncertainty increases, it becomes difficult to specify in advance not only the costs and benefits of regulation, but also its precise objectives. The government and its agencies will often have to go further and “negotiate” improvement targets with individual firms or clusters of firms. What is a good job, how many can be reasonably created, how technological and other firm-level choices influence job creation, what are the complementary policy levers that are available, how that set of instruments can be expanded – these are necessarily local, contextual questions. They can be answered, and periodically revised, only through a customized, iterative process of strategic interaction between public agencies and private firms. This process is alien to the familiar, principal-agent framework of rule-making which assumes that goals and social benefits must be known in advance if public action is to be effective and accountable. But it is the hallmark of the new type of regulation to which we turn next.

3. Key features of regulation under extreme uncertainty

Consider first contracting under uncertainty between private parties (which as we will see in a moment closely approximates the ARPA case). Under stable conditions each party can specify precisely what it expects in exchange with the other—*do ut des*. Precision, moreover, is often unnecessary, because in stable circumstances the same parties often contract repeatedly with each other; and these relations give rise to shared norms and expectations that guide performance even when there are gaps and ambiguities in formal agreements.

But under uncertainty the very trajectory of technology is unforeseeable and solutions in any domain are often found by applying ideas that arise far afield. It is neither possible to specify obligations in advance nor to rely on shared norms as supplements or substitutes for detailed agreements. Operating at the edge of established solutions, neither party can say exactly what is feasible, let alone what the other should contribute to the joint effort. When solutions are in view, they will often involve collaboration not with familiar partners, but with strangers, with norms and expectations of their own.

Under these circumstances the nature of contract itself changes. Instead of defining precisely each party’s obligations, the agreement establishes broad goals and a regime for evaluating achievement of them. As observed in domains as diverse as biotechnology, IT and advanced manufacturing this regime establishes regular, joint reviews of progress towards interim targets or milestones, procedures for deciding whether and with what exact aim to proceed or not, and mechanisms for resolving disagreements. The information exchanged under such a regime allows the parties to develop a more and more precise idea of the shared goal while allowing each to assess with increasingly reliability the capacities and good faith of the other: to observe if the capable stranger can become a reliable partner

and the long-trusted partner is capable of innovative tasks. As collaboration progresses, each party comes to rely increasingly on the capacities of the other, deterring opportunistic defection and generating or activating norms of reciprocity. Joint regular review and deliberate consideration of the interim results thus create the conditions in which informal norms and self-interested calculations bind the parties to continue promising collaboration in good faith. Trust and mutual reliance are the result of agreement to collaborate, not its precondition, just as the precise aims of cooperation are the outcome, not the starting point of joint efforts (Gilson, Sabel & Scott, 2009).

Regulation under extreme uncertainty arrives at a closely related solution from a somewhat different starting point. Under stable conditions mitigation of externalities is mandated by legislation and given precise form in consultation between the regulator and the regulated parties (subject to judicial review in case of continuing, insistent disagreement). The costs of mitigation are known to the regulated party but not (or at least not easily) to the regulator. Addressees of regulation try to use this information asymmetry to minimize their costs of adjustment while regulators devise ways of eliciting serviceable cost information without being captured by the actors that provide it. The upshot is a fixed set of limit values for emissions of pollutants of a particular kind and a schedule of fines for exceeding them.

Under uncertainty neither the regulator nor the regulated parties have reliable information on the possibilities and costs of adjustment in the medium term, and only vague conjectures regarding the possibilities that will open—or not—upon further investigation. Again the response—seen in food safety, civil aviation, and pharmaceuticals among many others—is the creation of an information exchange regime that ties ongoing specification of goals—here regulatory standards—to continuing exploration of new solutions. Typically the regulator, acting as before under a legislative mandate and after extensive consultations, establishes an ambitious, open-ended outcome: for example, “good water,” as measured by minimal deviation from the pristine state of a particular type of body of water such as an alpine stream or Mediterranean river; or a dramatic reduction, over an extended period, in vehicular emissions from various sources. The regulated entities—private parties, states or member states and their subdivisions in the US or the EU—are obligated to make plans to achieve the goals and to regularly report their results. Penalties in this regime are not calculated to deter infraction of clear rules but rather to incentivize cooperative production of the information from which standards will eventually be derived. Thus penalties are imposed as a rule only for failure to report or to report honestly, or for persistent failure to achieve results whose feasibility is demonstrated by the attainments of others in like positions; though infrequent, those penalties can be dauntingly severe, often amounting to exclusion from the market or (for public addressees) severe limits on decision-making autonomy. In contract law such information forcing sanctions are often called penalty defaults and we adopt that term here.

The combination of ambitious, open-ended goals, planning obligations and threat of potentially draconian penalties for obstinately uncooperative behavior encourage investigation of new possibilities, including contextualized variants of general solutions, and collaboration among regulated parties and between them as a group and the regulator. As long as some actors are looking to set new standards through their innovations—creating markets for innovative technology they develop, or simply putting competitors under pressure to match their performance—others will be less willing to cling to the status

quo at the risk of being caught out when methods advance. In an environment where the development of technology is uncertain precisely because of the continually surprising abundance of opportunities it affords, the expansive search for innovation is likely to feed on itself, with inquiry generating more inquiry, if only to minimize the chances of being surprised by developments. Search is likely to be collaborative either because projects are interdisciplinary and require the combined efforts of different specialists or because any one approach, interdisciplinary or not, is likely to fail and many actors will consider it prudent to pool the risks of exploration through various forms of collaboration.

Taken together many concurrent searches will yield a stream of surprises, unsettling understanding of what is technically possible and raising questions about what regulation can and should reasonably require. “Notice and comment”—the one-time consultation of stakeholders required in rule making by regulatory agencies in the US—gives way to regular, organized exchanges as regulators and addressees seek to establish common expectation in the face of rapidly evolving knowledge. Mutual ignorance and fear of surprises further bolsters information sharing between public and private actors just as it does among the latter. By making it risky to bet on the status quo and potentially rewarding to try to surpass it this regulatory regime turns uncertainty itself from an obstacle to demanding standards into a spur to collective learning that shows, cumulatively, how to realize them.

4. The environmental and R&D analogues

In this section, we discuss two successful examples of regulation—understood in the broad sense of public measures addressing externalities—under dynamic and uncertain conditions. The first is the Defense Advanced Research Projects Agency (DARPA) and its offspring, the Advanced Research Projects Agency-Energy (ARPA-E). They respond to the characteristic learning externalities that arise at the far frontier of science and technology, where for now there may well be no solution at all to a particular problem, and the search for one will likely end in costly disappointment. Worse still for the private investor, even when the search is successful it is unlikely that the daring pioneer can appropriate the returns from the discovery. The predictable result is underinvestment, from the standpoint of society as a whole, in research and technology. The second example is regulation by the EU and Ireland of Irish water quality and the Irish dairy industry generally. It illustrates the distinctive difficulties associated with mitigation of environmental externalities: Even when solutions can be developed in principle, it is difficult to estimate the costs of applying them, especially since, to be effective, general measures must be adapted to highly differentiated local circumstances. The familiar result is regulation that, for fear of imposing intolerable burdens on regulated parties, is often too timid to be effective; or when resolute, regulation that is ineffective for failure to take account of local particularity. Neither case is perfectly congruent to the “good jobs” challenge. But the success of both is due to the emergence of common mechanisms of governance under uncertainty that, we will argue, can make the good jobs strategy workable and accountable.

DARPA and ARPA-E

In discussions of industrial policy the DARPA, created in 1958 in response to the Soviet launch of the sputnik satellite, is often and usefully invoked as a reminder that the knowledge economy was not created solely by private actors—entrepreneurs, venture capitalists and technologists—responding only to opportunities signaled by markets. Far from being a mere bystander, the state, acting through DARPA and related agencies, played, and continues to play, a fundamental role in organizing the research from which are hewn the building blocks of the information economy. Among its iconic contributions are the computer network protocols underlying the internet, precursors to the global positioning systems and fundamental tools and devices for microprocessor design and fabrication. The accomplishments of DARPA have inspired a number of research agencies on similar lines, of which ARPA-E—a program created in the wake of the financial crisis to foster innovation in the energy sector—is both the most successful and the most faithful to the procedures of the original model.

Recent studies of ARPA-E examine in detail the institutional mechanisms by which such public entities can orient, coordinate and discipline collaborative investigation at the outer edge of technical possibility. If those mechanisms are today commonplace or rapidly becoming so it is not because DARPA's methods are widely emulated but rather because more and more organizations, public and private, are adapting to the high-uncertainty environment, once exotic, that shaped DARPA from the first.

At every stage in the organization of research—the definition of programs of investigation; selection of a portfolio of projects advancing the program purpose; and supervision of individual projects in the portfolio—ARPA-E treats goals as provisional, or corrigible in the light of experience. As with the contracts among innovating parties discussed above precise goals are the result of a search, not clear waypoints guiding it from the first. And at every stage in the process ARPA-E's program directors (PDs)—recruited from academia, industry or elsewhere in government for three-year terms, with the possibility of renewal—play a key role in the collaborative setting and revision of goals.

To begin with, PDs are hired largely on the basis of their promise in giving direction to an emergent area of investigation. For example a candidate with a background in geology will be hired to create a program in advanced geothermal energy. ARPA-E's overarching goal in establishing programs is to eliminate “white spaces” in the landscape of technical knowledge: missing capabilities, just beyond the frontier of current technical possibility, which, if mastered, would clear the way to advances in an important domain. A program might, for example, aim to support investigation of novel battery concepts with the potential to reduce storage costs by enough to make an attractive class electricity grid designs economically feasible.

Once program goals have been generally framed, the PD does a “deep dive,” supplementing and correcting her own background experience with reviews of the scientific literature and site visits to universities and companies by ARPA-E technical staff, as well as commissioned external studies and consultation with Department of Energy research managers. PDs then submit the emerging plan of

investigation to a first test of practicality by convening technical workshops involving leading engineering, scientific and commercial experts in well-defined domains. If the research plan, adjusted to reflect the exchanges at the workshop, passes review by all the PDs and the director of the Agency, a project is formally created as a component of the developing program, where a typical program (one of roughly 15 running concurrently) consists of 10 projects, each awarded \$3 million dollars to be spent over three years.

Proposals, many of them submitted by partnerships among several organizations, are developed and then executed in the same iterative manner, with goals open to recurrent challenge and revision. They begin with submission of a concept paper: a short document explaining why the proposal is superior to alternative approaches, and the applicants' response to foreseeable technical and commercial risks. External reviewers rate the proposals and their comments, along with the concept papers, pass to a committee, usually chaired by the PD, which selects proposals to be developed into full applications. The review process is repeated with the complete applications, except that applicants may rebut criticism by external reviewers, and their rebuttals are considered by the committee in the final selection of research partners. PDs and award recipients then have three months to agree on the precise milestones that will structure governance of the relation between ARPA-E and the research partner or "performer" as research proceeds.³

In the argot of ARPA-E this relationship is called "active project management," and it exhibits a strong family resemblance to the information-generating regime in contracting for innovation. Its most conspicuous feature is the quarterly progress report that performers must provide and that PDs and Agency staff rate with a traffic light system: red for projects that missed a critical milestone and are at risk of failing (which occurs in only 7 percent of project/quarters reviewed in a recent, comprehensive study); yellow for projects that missed a milestone but can be expected to recover (and stand very good chances of doing so according to this same study); and green for projects that are on track to reach their goals. Red ratings—or the anticipation of them—touch off an intensification of the routine site visits, conference calls, meetings, conference calls, and written analysis of problems and possible solutions by which PDs keep abreast of project developments. When projects struggle, calls and meetings can be scheduled weekly and milestones reset to permit an alternative to the failed approach. If recovery efforts fail, the PD sends an "at risk" letter warning of the possibility of termination, usually between one and two quarters before a project is in fact ended.

But collaboration between the research teams and the Agency, usually in the person of the PD, is more pervasive than the rating system alone might suggest. Milestones—which effectively mark the (changing) direction of the investigation—are added or deleted in fully 45 percent of the projects. And this number still significantly understates the extent of deliberate course correction because it does not

³ Formally cooperation between ARPA-E and recipients of research awards is governed by a cooperation agreement that specifies that the "Prime Recipient...is required to participate in periodic review meetings [to]...enable ARPA-E to assess the work performed under this Award and determine whether the Prime Recipient has timely achieved the technical milestones and deliverables" listed in an attachment. A sample agreement is available at <https://arpa-e.energy.gov/?q=site-page/funding-agreements>.

include modifications in the substance of milestones, which are said to be frequent, and which, as just noted, can amount to redirection of efforts, at least in the case of at risk projects. In addition project budgets can be increased (if interim results are especially promising) or decreased (when key milestones are persistently missed). In short the Agency rejects the model of hands-off, bet-on-the-person-not-the-project administration preferred by many established and successful research funders, public and private, in favor of the continuous, collaborative review and adjustment adopted in biotechnology, advanced manufacturing and venture capital.

ARPA-E is too new to permit any evaluation of its long-term impact: The energy industry—where even demonstration projects require substantial investment, innovators immediately confront legacy providers and regulation is more likely to constrain innovation than, as in pharmaceuticals, accommodate it—changes so gradually that large transformations only slowly become visible. But the available evidence does strongly suggest that ARPA-E is indeed choosing projects in the zone of uncertainty—where the positive externalities of research and development will be especially large—and using its information-generating regime effectively to make the most of its choices.

Expert disagreement about what is possible is a good working definition of uncertainty.⁴ If ARPA-E funds uncertain projects it should select projects whose prospects the best experts—its reviewers—disagree. This is what we observe. There is a very slight correlation between reviewers' ratings of projects and the likelihood that they will be funded. Selection is not based on a consensus view of project prospects. Perhaps more tellingly, holding the rating constant, the Agency picks the project where the range of reviewer rankings is the greatest—where judgements diverge the most. Plainly the PDs and the selection committee are relying on other information—rebuttals, observation of the research in workshop dialogue with peers, and much else besides.

The Agency's internal, traffic-light ratings of projects, and increase or decrease in resources dedicated to them as a result, are associated with the likelihood of achieving outward success: publishing a scientific paper, registering a patent, or securing some form of "market engagement," such as private financing, the formation of a company formation or release of commercial product. Budget increases, which of course reflect high ratings, increase the chances of market engagement for a project by 20 percent. Favorably rated project are indeed more likely to achieve all three outcomes, while projects with a high proportion of red quarters are significantly less likely to achieve any. Project selection and governance, moreover, do not seem to favor either scientifically oriented projects doing basic research validated in journal publications or commercially oriented projects doing applied research validated by patents or market engagement. Compared to projects in other branches of the DOE which are classified as doing either basic or applied research, ARPA-E projects have a higher rate of patenting and the same high rate of publishing. Most strikingly they are more likely than the specialized projects to produce both a publication and patent (Goldstein and Narayanamurti, 2018). A plausible interpretation is that they combine practical invention with scientific discovery on the model of use-inspired basic research made

⁴ Frank Knight classically distinguished risk—where an outcome is unknown but its probability can be estimated—from uncertainty—where it is impossible even to estimate the probability of future states of the world.

famous by Pasteur. As we will see next, commercial constraints and the penalty defaults imposed by EU environment law have made use-inspired research on similar lines central to the regulation of the water quality in Ireland and its dairy industry generally.

Irish dairy farming

Regulation, and especially environmental regulation, differs from ARPA-E's contractual governance of research in two ways. First, agreements between the Agency and award recipients are fully consensual (candidates compete for awards). Many addressees of regulation prefer no public constraints on their behavior; some actively resist imposition of rules. Penalty defaults therefore play an important role in inducing cooperation with the regulator, but none in the formation of award agreements. Second, ARPA-E faces uncertainty that arises from manifest limits of our knowledge of science and technology: the "white spaces" mark the places where we don't know the laws of nature that apply to a particular problem. Environmental regulation encounters such frontier uncertainty as well; but it is often challenged instead or in addition by uncertainties arising from the singularities of place: the way known factors—familiar pollutant streams; types of subsoil and geology—combine in particular contexts to produce unforeseeable results. "White spaces" get filled in once and for all. Once we learn the electrochemistry of cutting energy storage costs by a certain amount, *that* problem is solved. But environmental problems typically have to be re-defined and addressed place by place. In this regard environmental regulation strongly resembles, and can serve as a partial model for regulation of the good-jobs externality. In both a central task of governance is creating an information exchange regime that induces the local actors to cooperate to contextualize solutions while enabling them to benefit from the pooled experience of others, and vice versa.

Within environmental regulation non-point source pollution is the paradigmatic case of contextual uncertainty. The regular emissions of large polluters, such as power plants or sewage treatment facilities, are (relatively) easy to detect and control. Intermittent emissions from diffuse sources, such as the runoff from sporadic detergent use in scattered households, are not. Agricultural runoff is especially refractory because of the great variation in the pitch and absorptive capacity from field to field, the stark seasonal variations in weather and the rapid changes in the level and nature of productive activity induced by cycles of cultivation. We look to advances in the regulation of water pollution in agriculture to refine ideas about the governance of contextualization of the good jobs strategy, and to Ireland in particular, where pressures to reconcile demanding legal requirements to limit pollution with the needs of an expanding dairy industry have produced both an especially sharp understanding of the problem of contextual uncertainty and innovative reforms to address it.

The conviction that the dairying could be a modern engine of growth and the corresponding concern that environmental constraints could limit its expansion came late to Ireland. Through much of the 20th century Irish dairy farming, like Irish farming generally, was dominated by extremely small holdings, with limited export opportunities and relatively low productivity and incomes. Membership in the European Economic Community (the predecessor of the EU) and its Common Agricultural Policy (CAP) in 1973

expanded market access and raised prices, leading to increased output and productivity. The imposition of EU milk quotas in 1984 prompted consolidation in the sector, yielding a smaller but more efficient and capable cohort of specialized dairy farms still small—measured by farm acreage and herd size—in comparison to industrial producers. The Irish coops also consolidated and became first-tier suppliers of ingredients to global consumer food firms, which built processing plants in Ireland.⁵ In 2017 Ireland—which accounts for less than 1 percent of global milk output (Eurostat 2017 and FAO, 2018:5)—supplied almost 15 percent of the world’s infant formula market and was the second biggest exporter of infant formula to China.⁶ Altogether, Ireland exports 90 percent of its dairy output.⁷

Grass is the source of the competitiveness of Irish dairy. The larger representative Irish dairy farm has the lowest cash cost to output ratio of the key international milk producing regions, including the US, New Zealand and Australia (Thorne et al., 2017: 70). Home-grown, grass feed is much cheaper than purchased-feed concentrates; its price is relatively stable, while the price of purchased feed fluctuates with the price of the fuel used to produce it, sheltering Irish dairy farmers against a substantial risk. Cows that pasture on grass produce milk solids of superior quality; and the grazing cow is the emblem of food production at its most natural. Moreover, as the Irish quickly realized in the aftermath of the Great Recession, the share of imports is much lower in the inputs (grass first and foremost) of dairy firms than in the inputs of the transnational pharmaceutical and IT firms that dominated the Celtic Tiger boom in the years before it. The profits of domestically owned dairy firms remain in Ireland, while those of the high-tech firms are repatriated to their foreign owners. Per unit of output and exports the agri-food sector thus makes a larger contribution than the sectors dominated by FDI to the balance of payments and employment, as well to regional and rural development.⁸

For all these reasons the Irish dairy sector—farmers, farmer organizations, producer coops and their industry associations—together with its counterparts in various government departments have since the turn of this century come to see the national system of grass-based dairying on family farms as a model of production with a bright future, and a central role in the overall development of the country—provided it can reconcile increasing efficiency with regulatory and consumer demands for environmental sustainability.⁹ Wary engagement with climate change mitigation—not to say resistance—is giving way to active collaboration in developing new measures and institutions.

EU law compelled Ireland to respond, haltingly and reluctantly, to the pollution of water and air caused by agriculture long before farmers, farm organizations, dairy coops, the state extension service—

⁵ The consolidation of Irish milk processing is less pronounced than in Ireland’s major dairy export competitors, such as Denmark, the Netherlands and New Zealand, where one company processes as much as 70 or 80 percent of the milk pool (Prospectus, 2003).

⁶ <https://www.bordbia.ie/industry/manufacturers/insight/alerts/pages/chinadairyimports.aspx>

⁷ <http://www.bordbiavantage.ie/market-information/sector-overviews/dairy/>

⁸ It is estimated that in 2008 the wider agri-food (which includes beverages, infant formula and food ingredients) accounted for 40 percent of the total net foreign earnings of all primary and manufacturing industries (Riordan, 2012).

⁹ The Food Harvest 2020 strategy of 2010 set a target by 2020 of a 50 percent increase in the volume of milk production over the average of 2007-2009 milk supply (4.93 billion liters). The volume of milk production in 2017 had reached 7.27 billion liters, an increase of 47 percent (CSO 2018) - <https://www.cso.ie/en/releasesandpublications/er/ms/milkstatisticsdecember2017>.

Teagasc—became active advocates of sustainability. The Nitrates Directive of 1991 was one of the first measures to protect water quality from pollution by agricultural sources. It was highly prescriptive, setting out precise concentration limits that are transcribed in each member state’s Nitrates Action Program (NAP). Farms that fail to comply can be fined or disqualified from the EU single farm payment. Countries, such as Ireland, which fail to meet national limits must submit a plan for improvement to secure a temporary derogation of requirements or face the potential application of draconian sanctions typical of penalty defaults.

The Water Framework Directive (WFD) of 2000 has, in contrast, extremely broad objectives: “good water,” including minimal pollution by listed chemicals and “good ecological status,” defined for each type of water body (such as alpine streams or freshwater lakes) as minimal deviation from the distribution and quantity of life forms associated with a pristine or undisturbed body of water of that type. An “inter-calibration” procedure assures that countries apply comparable standards (Poikane et al., 2014). The basic unit of management is the river basin or catchment: the contiguous territory that drains into the sea at a single river mouth, estuary or delta. Member states produce a six-year River Basin Management Plan (RBMP) for each basin by a collaborative process in which public officials, experts and stakeholders specify objectives as well as procedures for translating them into concrete activities. Each member state appoints a water director to oversee execution of the plans. Together the water directors form a council which, in consultation with the EU’s executive body, the Commission, directs preparation of guidance—known as the Common Implementation Strategy (CIS)—in the application of the Directive. Until 2027 countries¹⁰ that fall short can submit a new RBMP at the end of each planning cycle on the grounds that the earlier approach proved technically infeasible, disproportionately expensive, or was obstructed by extraordinary natural conditions. Thereafter, as a penalty default, cost and feasibility will not excuse non-compliance.

Implementation of both Directives has proved frustratingly difficult. Adherence to “good practices” in agriculture, for instance, has often failed to produce improvements in nitrate levels; effective, inclusive participation of local actors in the definition and continuing revision of the intentionally open-ended goals has been a major stumbling block in the application of the WFD. The CIS has been revised many times¹¹; the WFD will be revised in 2019, among other reasons to reset the penalty default, as many member states will fail, even with energetic, good-faith efforts, to meet the 2027 deadline.¹²

¹⁰ For the guidance documents see http://ec.europa.eu/environment/water/water-framework/facts_figures/guidance_docs_en.htm. For detailed discussion of the CIS as an experimentalist institution at the heart of the WFD see Scott & Holder (2006).

¹¹ Hering, Daniel, et al., "The European Water Framework Directive at the age of 10: a critical review of the achievements with recommendations for the future." *Science of the total Environment* 408.19 (2010): 4007-4019; Voulvoulis, Nikolaos, Karl Dominic Arpon, and Theodoros Giakoumis. "The EU Water Framework Directive: From great expectations to problems with implementation." *Science of the Total Environment* 575 (2017): 358-366; Giakoumis, Theodoros, and Nikolaos Voulvoulis. "Progress with the WFD implementation in five European basins: Significant differences but similar problems." *European Journal of Environmental Sciences* 8 (2018): 44-50.

¹² On the possibilities for expanding the grounds for derogation beyond “natural conditions” while maintaining pressure to strive for compliance see Water Directors Meeting, “The Future of the Water Framework Directive

In Ireland these kinds of failure triggered a series of research programs under the Directives aimed at deepening understanding and control of pollution flows at the catchment and field levels. These programs (linked with similar ones on other member states) have in turn helped generate a web of institutions that is coming to function as an integrated system of local governance of water quality, greatly expanding public participation in environmental decision making in the process.

Teagasc established an Agricultural Catchments Programme (ACP) in 2008, in preparation for Irish's application for derogation of some requirements of the Nitrates Directive. Six catchment areas, differing in soil types, geology and types of farming, were selected to monitor and model the relations among farm management practices, the migration of nutrients from their source to various water receptors and the resulting changes in water quality. Some 300 farmers participated in the program, each supported in the development of a nutrient management plan by a Teagasc extension agent, who could in turn draw on the additional expertise of 15 researchers dedicated to the project.

The ACP's key finding is that variations in soil and subsoil types, and the underlying geology are in combination so influential in the absorption and drainage of nutrients that general rules of nutrient management, let alone plans based on them, will regularly fail. For example, poorly drained fields with phosphorus values so low as to obstruct cultivation and seem environmentally innocuous may still pollute because of fast surface runoff. Conversely, soils with phosphorus concentrations in excess of agricultural needs may not pollute at all because they are especially well drained (Shortle & Jordan, 2017: 17). The policy implication is that a nutrient management plan should be a starting point or provisional guide for investigation, farm by farm, of environmental risks, and how most economically to address them. It would be only a slight exaggeration to say that the implementation *is* the plan, and it is co-produced by the advisor and farmer, collaborating in identifying the problems of particular farms, devising remedies, and jointly monitoring the results (Burgess, phone interview, October, 2018). A second catchment study, undertaken by the EPA, confirmed the findings of the ACP, extended investigation of the mechanisms of pollution transport to the geological structures beneath the soil layers, and showed that disruption of pollution pathways is often a more effective means of mitigation than attempting to eliminate pollution at its source or contain its effects at the receptor.

The new catchment program is part of a larger effort by the EPA and its partner institutions in water quality management to establish a cascading process of national, regional and local consultation to select priority areas for intervention and create local governance institutions to support execution of the agreed interventions with full and effective participation of the affected actors. The selection process and new governance institutions come together when the priority areas, once selected, are subjected to a final, searching review in "local catchment assessments": field-level examinations by the local actors themselves of the source of pollution in given water bodies. This assessment determines the local work plan, specifying, costing and prioritizing projects. Such collaborative review of priorities is particularly important in rural areas, where multiple, potentially significant small point pressures and diffuse sources of pollution (Daly et al., 2016: 161) can confound mitigation, and deep local knowledge is indispensable to a deliberate and consensual choice of which problems to attack and in what order. A new Local

(WFD) – Water Directors input to the fitness check process on experiences and challenges of WFD's implementation and options for the way forward," November 15, 2018.

Authority Water and Communities Programme (LAWPRO), a shared service between all local authorities, supports the catchment assessment by providing technical assistance to all the stakeholders (and indirectly by helping each locale develop the structures needed for inclusive engagement with the articulation and implementation of the RBMP).¹³ Agricultural problems detected by field assessments are referred to specialist Agricultural Sustainability Advisors who, in consultation with the assessment team, assist the implicated farmers to improve their land, farmyard and nutrient management practices.¹⁴ The corps of sustainability advisors link contextualization of water management at the catchment or territorial level to contextualization of pollution mitigation measures on the farm, completing the nascent system of local governance.

We draw three lessons for the design of the good jobs strategy from Irish and EU experience with environmental regulation of contextual uncertainty. First, while framework legislation (the WFD) and penalty defaults orient and incentivize the creation of new governance instruments for the local contextualization of general policies, making those institutions work in the contexts to which they apply requires continuing revision of initial plans of light of—frequently disappointing—experience. The recent flurry of institution building in Irish water regulation, the culmination of systematic investigation and hard experience, was preceded by many false starts and misdirected half measures. There are principles of design for these institutions, but no blueprints.

Chief among these principles—the second lesson—is that contextualization in the sense of recognition of the need for local solutions to idiosyncratic local problems is a corrective and supplement to higher-level decision making and procedures, but not a substitute for them. LAWCO review modifies the specification of local targets identified by national and regional review, and how and in what order they are approached. Local authorities and stakeholders are not free to disregard the national list or to insist that their interventions, whatever the outcome, are unquestionable solely because they have been locally authorized; rather, lower levels correct higher levels and vice versa. More generally contextualization is just a variant of the reciprocal review that we encountered in contracting for innovation and which typifies collaboration under uncertainty, outside the principal-agent framework.

Finally, contextualization blurs the distinction between regulation, directed to compliance with rules—order maintenance within a given system—and the creation of new institutional systems. Contextualization induces collaboration between regulators and other public officials and regulated entities in the development of novel forms of capacity building and novel forms of public participation in regulatory decision making. Irish dairy farmers in the catchment projects prepare their nutrient management plans with the support of specialist extension agents, who consult themselves with catchment specialists; farmers with environmental problems respond in collaboration with newly formed catchment assessment teams and a new corps of specialist sustainability advisors. Traditional extension agents propagate consolidated expertise. In co-developing improvement plans with individual farmers and each other these new specialists are reconsidering and revising current understandings as

¹³ For a good overview of the role of the Waters Catchment Assessment Team and its place in the new governance structure see <https://www.catchments.ie/the-local-authority-waters-programme-catchment-assessment-team/>.

¹⁴ On ASSAP see <https://www.teagasc.ie/environment/water-quality/farming-for-water-quality---assap/>

much as applying them. Collaborative investigation is necessary precisely because current rules and best practices run out; and establishing what should be done goes hand in hand with understanding and building the capacity needed to do it. The creation of LAWCO, with its open-ended implications for local governance, suggests how regulation under uncertainty can, in a further step, prompt institutional innovations that implicate basic ideas of democratic participation and administrative due process. When we speak, as we do next, of regulation in relation to the good jobs strategy we mean the term in this enlarged sense of fixing (and revising) requirements and inducing the creation of novel institutions, with all their further spillovers, that enable the addressees of regulation to meet them.

5. Applying the model to good jobs

The concept of “good job,” like clean water, is imprecise and needs to be operationalized in a way that is both evolving and context-dependent. Reasonable, attainable targets for good-job creation must remain provisional, to be revised under new information. We can think of them as rebuttable presumptions, mandating behavior except when there is compelling evidence that they demand the impossible, or do not demand enough. Achieving the targets depend on decisions on investment, technological choice, and business organization, the consequences of which are unknowable *ex ante*. Governance under uncertainty takes as its starting point the provisionality of ends and means and the need for disciplined review and revision. Here we sketch what the model would look like when applied to the challenge of creating good jobs through public-private collaboration. We stress similarities, but also some differences.

Simplifying greatly, the uncertainty government agencies face in the ARPA-E case is principally about technological feasibility. The uncertainty in the dairy case derives largely from local adaptation. A good-jobs program faces uncertainties of both kinds. Creating or preserving good jobs in a particular place often depends partly on extending technological capabilities: mastering techniques that are wholly novel (at least in some particular application) or so new to a given locale that they must almost be reinvented to be mastered. Here ARPA-E’s experience with active project management and collaborative review and adjustment of milestones is directly relevant. But fostering good jobs depends at least as much on solving highly idiosyncratic, place-specific problems: failures of coordination between local firms and training institutions; between firms and their (potential) supply-chain partners; and the managerial breakdowns or skill gaps within individual firms and institutions to which the coordination problems point. Here the peer assessment of local problems and new forms of collaboration with networks of extension experts developed in Irish pollution control come into their own. There are many ways to imagine integrating, or coordinating the operation of the two variants of the governance model in particular conditions; how precisely is a practical question, to be answered in context when the time comes, and provisionally—subject to correction—in accord with the precepts of the governance model itself.

An immediate question has to do with the timing, scale and scope of the obligations (and penalty defaults) to be imposed on private firms. If there is a genuine good-jobs externality, and a national or

subnational mandate to address it, there is no reason why in principle the obligation to do so should not be applied immediately to all firms in the relevant jurisdiction. But as just noted those obligations are inherently broad, open-ended and, at least initially, ill-defined. They would begin with the requirement to make plans to progress towards forms of organization and deployment of technology that in combination produce better jobs, and to make such plans in coordination with relevant peers and institutional partners. But this may well be a too draconian first step. Unless we assume extraordinary consensus in favor of addressing the jobs externality or a dangerously coercive state authority, we can't really imagine the regulator imposing on all firms, or even all firms in certain sectors, the obligation to make such plans; and if can't imagine that still less can we envisage penalty defaults for persistent failure to make good faith efforts to comply.

It is easier to imagine imposing such requirement and penalties on actors who volunteer to participate in government programs designed to achieve the same outcome, and conferring benefits in the form of improved regulation, better coordination, extensive customized support services, or the like in return for participation. The framework goals, continuous monitoring and reporting requirements and penalty defaults (in the form of exclusion) would apply in this setting; but they would be the mutually agreed, common governance mechanism of a whole portfolio of industrial policy measures addressing the good-jobs externality. The voluntary and selective nature of the partnership with state agencies suggests that this start-up phase the good jobs strategy could make use of ARPA-E's governance of program definition—proceeding incrementally, and repeatedly exposing designs to objections and alternatives—and active project monitoring.

A key benefit of these voluntary arrangements over the medium term is to develop an inventory of “good practices”—a repertoire of contextualization measures variously suited to a wide range of settings—that can eventually guide application of the good jobs strategy to a larger set of firms, cutting the costs and increasingly the chances for early successes of broader coverage. Put differently the initial, selective projects would serve as a pilot program for the new system of regulation, with the qualification that pilots are usually understood as practical tests of promising concepts, whereas in this case their purpose would be more to identify and begin to refine promising approaches under real world conditions than subject them to definitive tests. As formal obligations are extended, the arrangements would come to resemble the European regulatory model, with a uniform requirement of participation but responses highly differentiated by locale. The need for contextualized support for the less capable actors drawn into the system would grow apace.

An intermediate arrangement might also be possible. Firms might be asked to make a choice between participation in customized good-jobs compacts with public agencies and submitting to a fixed regulatory regime that imposed a common, universal set of benefits and obligations linked to job creation (for example, a schedule of tax incentives/penalties in return for an increase of $x\%$ per annum in number of employees at wages at or above $y\%$ of the local median wage, where x may be tied to the state of the business cycle). Firms would then self-select into their preferred regime. The choices firm make would be informative about the relative effectiveness of the two regimes, feeding in turn into their revision over time.

With these design principles and staging practicalities in mind a good-jobs industrial policy could be introduced in four steps. First the government commits in legislation or by other means to address the problem of bad jobs and no jobs as a constitutional externality that threatens the foundations of our democracy and requires for its solution concerted cooperation between regulators, service provide and private actors. The framing legislation mandates regulators with relevant authority to put in place information-generating regimes that allow for standard setting and revision. The same legislation creates an inter-agency body to periodically review, and prompt improvement of regulatory responses, and resolve coordination problem arising from them, and provides funds and authority for voluntary programs in anticipation of an eventual, step-wise extension of regulatory reach.

Regulators who currently have delegated authority for areas directly affecting job abundance and quality—vocation training, agricultural and manufacturing extension, standard setting and the like—introduce, in a second step, innovation-inducing and contextualizing governance mechanisms where these are not already in place, anticipating the need for support services to help vulnerable actors comply with increasingly demanding requirements. The requirements can take different forms, including specific employment quantity targets and/or standards.

Where current regulatory authority doesn't reach, the government creates volunteer, public private programs to advance the frontiers of technology and organization, or—and of equal and perhaps greater importance—provide support services and perhaps subsidies to help firms bridge the gap between their current low-productivity/low skill position and participation in the advanced sector. These programs, in their ensemble, would have to combine services to workers as well as managers; they would have to be customized to the needs of particular sectors and locales, and probably both. They would adhere to the design principles of innovation-inducing governance; their performance would be accordingly reviewed, and their goals adjusted, by the responsible agency, and then, if problems persist, the inter-agency body.

Finally, conditional on the success of voluntary arrangements, the scope of these practices would gradually be made obligatory for non-participating firms, starting requirements for submitting credible plans for the improving the quality and quantity of jobs together with competitive position by better organization, use of skill and technology, where appropriate in coordination with other firms and institutions. Penalty defaults would be imposed on laggards firms that, despite the availability of support services, persistently fail to comply.

To place our proposed framework in sharper relief, we discuss briefly how it relates to some existing initiatives for promoting manufacturing and job creation.

Comparison with current initiatives

Of the three major components of the building good jobs program—extension services and cooperative research programs for existing firms; job creation and attraction policies; and active labor market

policies or workforce development—it is to the last, workforce development, that the governance practices we described previously in advanced technology and European regulation have been most consistently applied to adapt intervention to local contexts—with demonstrable success, inducing emulation. In the case of job attraction the experience has been nearly the reverse: Local and state politicians outbid each other to win outside investments in new facilities, more often than not in deals that (in contrast to the contracts for innovation under discussion here) specify all terms of the exchange fully in advance: so much in subsidies for each job created at an agreed wage rate and in a facility of an agreed type. Though there are important exceptions, in conspicuous, recent cases local learning is scarcely an afterthought. Extension services and cooperative research programs are an intermediate case, with successive waves of institutional innovation leaving the policy landscape dotted with small organizations that appear to do some good in their ambits, but in their isolation do not much effect the course of development (Block et al., 2018; Deloitte, 2017). These different outcomes, so far as we can see, reflect the vagaries of policy choices and economic flux, not the inherent ease or difficulty of pursuing contextualization strategies in the various domains. To all appearances, in fact, workforce development should be the most refractory terrain, since programs must engage at-risk groups and address many of the compound problems—financial, educational, familial—that notoriously vex social welfare services. The focus here therefore is on workforce development to illustrate an important application of our general governance principles to building good jobs; we refer to some prominent, recent cases of job attraction to underscore the difference between an approach which is deliberately sensitive to the uncertainties of context and one which deliberately is not.

Many of the most successful workforce development programs trace back to or were in part inspired by Project QUEST (Quality Employment through Skills Training), founded in San Antonio in 1992.¹⁵ In the late 1980s San Antonio was hit by a wave of plant closings—an early portent of broader dislocations to come. The displaced workers lacked the skills for the new jobs being created in health care, IT and other sectors; the service sector jobs for which they were qualified paid too little to support a middle-class family. Two faith-based social movement organizations (affiliates of the Alinsky Industrial Areas Foundation), seeing the urgent need for a program to equip the region’s largely Hispanic population for good jobs, secured municipal funding to create Project QUEST (Warren, 2011).

As a pioneer in this area Project QUEST faced a double challenge: On the one hand it had to identify emerging opportunities on the local labor market, alert the city’s community college system (then still inattentive to business needs) to them, and help shape the substance and timing of new course offerings to meet the requirements of both firms and students. On the other hand it had to learn to support a population of learners at high risk of failing at school: learners almost of whom needed to pass difficult remedial courses in order to qualify for further study, and many of whom were older, with family and financial burdens on top of anxieties about returning to school.

From the start the social movements backing Project QUEST were determined to avoid repeating the mistakes of the failed job training programs they could see at hand. In facing their challenges, therefore,

¹⁵ Unless otherwise indicated this account of Project QUEST draws on (Rademacher et al., 2001).

they turned to former military men with long experience in workforce development. The first executive director was the former commander of the Air Force Recruiting Service; his successor, and many managers they hired, had a similar background. These managers brought with them not respect for military discipline and hierarchy but rather for the culture of continuous improvement—the continuous monitoring of individual cases and rapid learning from disruptions at the core of our governance principles—that took root in many parts of the US military before it became standard operating procedure in much of the economy.¹⁶ An expression of this culture was the early creation of a dedicated management information system, highly unusual for an organization like Project QUEST at the time, to track the performance of individual students, both to keep their counselors abreast of their progress and to allow continuing review of overall organizational performance.

To be eligible to participate in QUEST students must demonstrate need (generally earnings of less than 50 percent of the local median wage) and levels of literacy and numeracy sufficient to ensure reasonable chances of succeeding at the intense remedial programs typically needed to prepare for the required, basic courses in community college programs. Once admitted students design in collaboration with a counselor an Individualized Service Strategy (ISS): a bundle of "wrap-around" services and supports to help them get through training. The ISS includes financial support for tuition and other school-related expenses, but extends, according to need (and with the help of outside partners), to subsidies for child care and rent. In particular cases the ISS will address problems of transportation, health or domestic violence—in short, any stumbling block on the path to completing training.

Counseling is continuous and intense. Students meet their counselors individually, but also in small, stable groups in weekly, hour-long Vision, Initiative and Perseverance (VIP) sessions, where they share problems and devise mutual support strategies. A key purpose of these meetings is to identify and respond to emergent problems before they trigger a cascade of failure ending in withdrawal from the program. In effect the counseling sessions in combination with information of students' class performance allows for continuous adjustment of the ISS.

A recent RCT evaluation of the earnings of QUEST participants nine years after leaving the program demonstrates the effectiveness of providing a bundle of individualized—wrap-around—services and supports (Roder and Elliott, 2019). QUEST participants earn roughly ten percent more per year than the control group; and while the gap varies from year to year it does not diminish—and may be growing—over time. Crucially the difference in earnings is the greatest for the most at-risk subgroups: students who took part in QUEST when they were older than the normal school-going age, with children, and so on. The treatment worked best for those who most needed it—an unusual finding in the literature evaluating the effects of social welfare interventions.

In recent years, as community colleges are increasingly drawn into training partnerships with local firms, more students from more diverse backgrounds seek new qualifications, and the shortcomings of limited, "light touch" interventions to increase completion rates become more conspicuous, the schools

¹⁶ For the transformative effect of the introduction of a demanding variant of continuous improvement developed in the US nuclear navy on the regulation of nuclear power generation see Rees, 1996.

themselves are successfully providing many of the individualized services originally offered by QUEST. A leading example is the Accelerated Study in Associate Programs (ASAP) of the City University of New York (CUNY).

Like Project QUEST ASAP provides financial support (to bridge the gap between the available aid and tuition and other fees) and wrap-around services, above all a dedicated advisor for each student who furnishes frequent and comprehensive support. Again a principle goal is to identify and resolve issues before a student drops out of school. This individual consultation is in high demand. An RCT evaluation of ASAP found that program group students reported an average of 21 counseling contacts per semester in their first year, while control group students reported an average of 4. In addition, to foster the kind of camaraderie and mutual support developed in VIP sessions ASAP reserves seats in courses for program students and schedules some classes back-to-back to mesh with their schedules. Program directors and a central evaluation staff routinely review student data and outcome trends to identify areas for improvement (Weiss et al., forthcoming).

Cumulatively, over the course of the three-year program, these interventions have produced a striking increase in completion rates. Nearly 40 percent of the students in the program group in the RCT study graduated by the end of the program; the graduation rate for control group members was 22 percent. This is largest effect estimate on three-year graduation rates, according to the authors of the study, of any large-scale RCT in higher education to date. What makes ASAP uniquely successful, they find, “is that its multiple, integrated, and well-implemented services address multiple prevalent barriers to student success, and those services are offered for three full years” (Weiss et al., forthcoming, np): in our terms, continuing contextualization based on continuous monitoring.

We would expect a program designed to adjust to local circumstance to be scalable, and ASAP is proving to be. Since 2014 three Ohio community college have implemented ASAP, and early impact assessments show results comparable to those obtained in New York (Sommo et al., 2018). Community college leaders are following the success of ASAP closely and devising their own systems of comprehensive support, with the goal, in the words of one, of “making help unavoidable.” In the most ambitious cases growing confidence in the ability to train low-skill workers is encouraging community colleges to enter extensive partnerships with large firms such as Amazon, where the school offers customized training and student support and the company pays tuition expenses, synchronizes demands on the time of participating employees to mesh with school needs and, perhaps most crucially, provides a career ladder from unskilled work into management.¹⁷

If workforce development programs are increasingly aware of the need for continuous learning in response to the uncertainties of context, programs that directly target employment creation by attracting inward investment seldom are. The most visible of such programs, typically administered by states rather than the Federal government, are tax incentives provided to large investors in return for specific commitments on job creation. The Foxconn and Amazon deals, in Wisconsin and New York,

¹⁷ Interview with David T. Harrison, President, Columbus State Community College, Columbus, Ohio, April 18, 2019.

respectively, are recent high-profile examples. The Taiwanese company Foxconn had agreed to create 13,000 well-paying jobs in Wisconsin in return for more than \$4.5 billion in government incentives. Amazon promised creating 25,000 jobs over a decade in return from an incentive package from New York valued at nearly \$3 billion. Both arrangements have blown up in controversial ways, and their failures are instructive in ways that demonstrate the superiority of the alternative approach we are suggesting here.

Essentially, the Foxconn and Amazon deals – as well as similar tax incentive programs – were predicated on ex-ante contractibility (and hence a stable environment). With enough predictability about market and technology conditions, firms can make rational calculations about employment commitments. And the states have the assurance that firms will deliver. Once the contract is written down, the state remains at arms' length from the firm. In Amazon's case, the company said it wanted cities to "think big." In reality, as one commentator has noted, "the creative thinking was exclusively focused on incentive offers" (Jensen 2019). There was no provision for an ongoing relationship to build trust, exchange information, and set and revise mutual goals. In such cases, the contract is either complied with or violated. If the firm turns out to be unwilling or unable to carry out the terms of the contract – as was the case with Foxconn and Amazon, the former because unforeseen changes in demand and technology and the latter because of unexpected political fallout – there is little room for revision or renegotiation.

Bartik (2018) has studied such tax incentive programs more broadly and concludes that, even when they work, are not very cost effective. This is especially true when local incentives have to be financed by cuts in public expenditures elsewhere (e.g., education or infrastructure). Bartik argues that the most effective employment programs are those that focus specifically on local labor demand and supply conditions. He emphasizes three strategies in particular. In the first bucket is the provision of customized public services to small and medium-sized enterprises. These include job training tailored to local employers and run by local community colleges, and "manufacturing extension services" that provide marketing and technology advice. The second bucket contains targeted investments in workers' skills and training, ranging from pre-school programs to wage subsidies. Third, he mentions infrastructure programs that increase land supply and thereby lower business costs.

These findings are consistent with our emphasis on iterative fine-tuning and evolving standard setting in lieu of ex-ante rules. The design of locally effective incentive packages along these lines obviously requires significant amount of information discovery and trial-and-error on the part of local development agencies. Therefore organizational arrangements of the type we have discussed here would seem to be crucial. Note also that while Bartik's (2018) focus is on manufacturing employment, our proposals would apply to service sectors as well. This is important since it is unlikely that the long-term, secular decline in the share of manufacturing employment can be reversed.

More broadly, good practice in industrial policy has moved away from presumptive approaches that assume the government has a good fix on the underlying problem and the requisite solutions. For example, industrial parks or enterprise zones presume the absence of good jobs is due to, say, high taxes and poor infrastructure and create spaces where neither is a problem. Such pre-packaged solutions

work poorly when firms face differentiated obstacles – lack of workers with appropriate skills or inadequate access to specialized technologies, for example. The collaborative framework we have outlined here has the advantage that it is explicitly diagnostic – i.e., focused on information discovery.

6. A reason not to despair and a reason to hope

We conclude on two positive notes. First, we argue that the prevailing academic pessimism about job creating strategies may be misplaced as it is based on conflating treatments with meta-treatments. Second, we note the possibility that the governance arrangements we have sketched out may also help enlarge the constituency for acting on the problems they address.

Consider first the point about policy pessimism. The literature we have referred to briefly in the introduction suggests there are few, if any, policies that work reliably to expand good jobs. The conclusion to which such findings often lead is either some combination of agnosticism (as in Austen et al., 2018) or a call for more randomized evaluations (as in Autor et al., 2019). Under our approach, these mixed results are not a surprise. The dual challenge of dealing with uncertainty and contextualization implies there are no fixed, clear-cut remedies. What is important is to get the governance regime right. With the appropriate regime in place, the hope would be that each locality can develop its own set of evolving practices. In the language of RCTs, our treatment is really a meta-treatment: a protocol for figuring out the treatment to apply in a particular setting. Correspondingly, proper evaluations would have to be carried out at the level of the governing regimes rather than individual policies.

When accurately presented, meta-treatments should have external validity. Of course, this is not without practical difficulty. If a meta-treatment is dressed up as a simple treatment, it is likely to have no external validity because each follower will interpret the costume in her own way – unless, exceptionally, she sees through the disguise to the governance it conceals.

Next, the point about building constituency. In many discussions of industrial strategy the principle problem is creating a coalition of public and private interests in favor of development. The first task is rallying a national or local coalition of private and public actors in favor of a growth strategy with clear, immediate objectives. The formation of consensus and the clarification of objectives leads naturally to the creation of public-private partnerships to advance particular projects. The public actors contribute their expertise and authority in, say regulation; private actors make complementary contributions with respect to markets and their firms. So long as the state retains sufficient autonomy to avoid capture the governance of the industrial policy projects is part and parcel of the consensus that underpins the public-private partnerships.

We agree that building political will is a threshold condition for industrial policy. Full deployment of the good jobs strategy would eventually require national mobilization. But under current conditions—when development is as likely to depend on exploring and building domestic capacities as accumulating know-how and capital; when uncertainty makes the selection of goals necessarily provisional and the revision

of ends and means routine—governance of a good jobs strategy can presume only a thin, initial, background consensus, and does not grow directly from it. On the contrary: Fear that, under uncertainty, ambitious and urgently needed programs cannot be effectively and accountably administered could cast a shadow over consensus building, causing some potential members of a coalition in favor of a good jobs strategy to back away from a risky venture.

An attractive feature of the approach here is that the same institutions of interactive governance that enable the parties to specify and solve the problems they face under uncertainty also enable them to develop the trust and mutual reliance they need to deepen and broaden their efforts. The broad coalition needed for the good-jobs strategy to succeed need not pre-exist; it can and will likely be the result of pursuing the strategy. Innovative modes of governance allow the parties, beginning with only a vague understanding of the substance and scope of their goals, to assess one another's capacities and good faith in the very process of refining ideas of what the eventual project should be. Trust and coalition building—the acceptance of mutual vulnerability—are as much or more the outcome of joint problem solving as its precondition.

Bibliography

Austin, Benjamin, Edward Glaeser, and Lawrence Summers, "Jobs for the Heartland: Place-Based Policies in 21st-Century America," *Brookings Papers on Economic Activity*, Spring 2018.

Autor, David, David Dorn, and Gordon Hanson, "When Work Disappears: Manufacturing Decline and the Falling Marriage-Market Value of Men," December 2018, forthcoming *American Economic Review: Insights*.

Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi, "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure," December 2017.

Autor, Davis, Aaron Li, and Matthew Notowidigdo, "Preparing for the Work of the Future," Abdul-Latif Jameel Poverty Action Lab (J-PAL), 2019.

<https://www.povertyactionlab.org/sites/default/files/documents/work-of-the-future-literature-review-4.2.19.pdf>.

Azoulay, Pierre, Erica Fuchs, Anna Goldstein, and Michael Kearney. "Funding Breakthrough Research: Promises and Challenges of the 'ARPA Model.'" NBER Working Paper 24674. Cambridge, MA: National Bureau of Economic Research, June 2018. <https://doi.org/10.3386/w24674>.

Ballard-Rosa, Cameron, Amalie Jensen, and Kenneth F. Scheve, "Economic Decline, Social Identity, and Authoritarian Values in the United States," October 2018.

Bartik, Timothy J., "What Works to Help Manufacturing-Intensive Local Economies?" Upjohn Institute Technical Report No. 18-035. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 2018.

Berk, Gerald. "Building the Problem-Solving State: Bridging Networks and Experiments in the US Advisory Specialist Group in World War II." *Politics & Society* 46, no. 2 (June 2018): 265–94. <https://doi.org/10.1177/0032329218773711>.

Block, Fred L, and Matthew R Keller. *State of Innovation: The U.S. Government's Role in Technology Development*. London: Routledge, 2011.

Block, Fred, Matthew R. Keller, and Marian Negoita, "Network Failure and the Evolution of the U.S. Innovation System," unpublished paper, 2018.

Colantone, Italo, and Piero Stanig, "Global Competition and Brexit," unpublished paper, Bocconi University, September 28, 2016.

Colantone, Italo, and Piero Stanig, "The Trade Origins of Economic Nationalism: Import Competition and Voting Behavior in Western Europe," BAFI CAREFIN Centre Research Paper Series No. 2017- 49, January 2017.

Colantone, Italo, and Piero Stanig, "The Economic Determinants of the 'Cultural Backlash': Globalization and Attitudes in Western Europe," August 4, 2018.

Dal Bò, Ernesto, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickn, "Economic Losers and Political Winners: Sweden's Radical Right," August 2018.

Deloitte, Manufacturing USA: A Third-Party Evaluation of Program Design and Progress, January 2017.

Frieden, Jeffrey, *Global Capitalism: Its Fall and Rise in the Twentieth Century*, W. W. Norton, New York, 2006.

Friedman, Benjamin M., *The Moral Consequences of Economic Growth*, Alfred A. Knopf, New York, 2005.

Gilson, Ronald J., Charles F. Sabel, and Robert E. Scott. "Contracting for Innovation: Vertical Disintegration and Interfirm Collaboration." *Columbia Law Review*, 2009, 431–502.
<https://doi.org/10.2139/ssrn.1289428>.

Goldstein, Anna, and Michael Kearney. "Know When to Fold 'Em: An Empirical Description of Risk Management in Public Research Funding." Working paper version: September 20, 2018. NBER, 2018.
<https://www.ssrn.com/abstract=3252548>.

Goldstein, Anna P., and Venkatesh Narayanamurti. "Simultaneous Pursuit of Discovery and Invention in the US Department of Energy." *Research Policy* 47, no. 8 (October 2018): 1505–12.
<https://doi.org/10.1016/j.respol.2018.05.005>.

Guiso, Luigi, Helios Herrera, Massimo Morelli, and Tommaso Sonno, "Demand and Supply of Populism," CEPR Discussion Paper DP11871, February 2017.

Jensen, Nathan M., "Five economic development takeaways from the Amazon HQ2 bids," Brookings, March 4, 2019 <https://www.brookings.edu/research/five-economic-development-takeaways-from-the-amazon-hq2-bids/>.

Khosla, Pradeep K., Paul T. Beaton, Committee on Evaluation of the Advanced Research Projects Agency-Energy (ARPA-E), Board on Science, Technology, and Economic Policy, Policy and Global Affairs, Board on Energy and Environmental Systems, and Division on Engineering and Physical Sciences, eds. *An*

Assessment of ARPA-E. Washington, D.C.: National Academies Press, 2017.

<https://doi.org/10.17226/24778>.

Mazzucato, Mariana. *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. Revised edition. Anthem Frontiers of Global Political Economy. London ; New York: Anthem Press, 2014.

Miller-Adams, Michelle, Brad Hershbein, Timothy J. Bartik, Bridget Timmeney, Amy Meyers, and Lee Adams, "Building Shared Prosperity: How Communities Can Create Good Jobs for All," Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 2019. <https://research.upjohn.org/reports/235>.

Osterman, Paul, "In Search of the High Road: Meaning and Evidence," *ILR Review*, 71(1), January 2018, 3-34.

Rademacher, Ida, Marshall Bear, Maureen Conway, "Project QUEST: a case study of a sectoral employment development approach, Sectoral Employment Development Learning Project case studies series," Economic Opportunities Program, Aspen Institute, Washington, D.C., 2001.

Rees, Joseph V. *Hostages of each other: The transformation of nuclear safety since Three Mile Island*. University of Chicago Press, 1996.

Remes, Jaana, James Manyika, Jacques Bughin, Jonathan Woetzel, Jan Mischke, and Mekala Krishnan, *Solving the Productivity Puzzle*, McKinsey Global Institute, February 2018.

Roder, Anne, Mark Elliott, *Nine Year Gains: Project QUEST's Continuing Impact*, Economic Mobility Corporation, New York, 2019.

Sabel, Charles, and Gary Herrigel. "Collaborative Innovation in the Norwegian Oil and Gas Industry: Surprise or Sign of a New Economy-Wide Paradigm?" In *Petroleum Industry Transformations: Lessons from Norway and Beyond*, edited by Taran Thune, Ole Andreas Engen, and Olav Wicken, 231-48. Abingdon, Oxon & New York, NY: Routledge, 2018. <https://doi.org/10.4324/9781315142456>.

Sommo, Colleen, Dan Cullinan, Michelle S. Manno, *Doubling Graduation Rates in a New State: Two-Year Findings from the ASAP Ohio Demonstration*, MDRC, 2018.

Sabel, Charles, Gary Herrigel, and Peer Hull Kristensen. "Regulation under Uncertainty: The Coevolution of Industry and Regulation." *Regulation & Governance* 12, no. 3 (September 2018): 371-94.

<https://doi.org/10.1111/rego.12146>.

Temin, Peter, *The Vanishing Middle Class: Prejudice and Power in a Dual Economy*, MIT Press, Cambridge, MA, 2017.

Warren, Mark R., *Dry Bones Rattling: Community Building to Revitalize American Democracy*, Princeton University Press, Princeton, 2011.

Weiss, Michael J., Alyssa Ratledge, Colleen Commo, Himani Gupta, "Supporting Community College Students from Start to Degree Completion: Long-Term Evidence from a Randomized Trial of CUNY's ASAP," Forthcoming, *American Economic Journal: Applied Economics*.

Weitzman, Martin L., "Prices vs. Quantities," *The Review of Economic Studies*, 41(4), October 1974, pp. 477-491.

Wilson, William Julius, *When Work Disappears: The World of the New Urban Poor*, Alfred A. Knopf, New York, 1996.