Economics & Philosophy (2024), 1–17
doi:10.1017/S0266267124000282

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

The body politic has private parts: market creation as a policymaking tool

Kirun Sankaran

Political Economy Project, Dartmouth College, 19 College St, 6035 Thornton Hall, Hanover, NH 03755, USA
Email: kirun.sankaran@dartmouth.edu

(Received 05 May 2024; revised 12 July 2024; accepted 26 July 2024)

Abstract
Philosophical arguments about government contracting either categorically oppose it on legitimacy grounds or see it as largely anodyne. I argue for a normatively distinct kind of contracting – the advance market commitment, or AMC – and show that it is justified by the same liberal values that justify the welfare state.

Keywords: Contracting; privatization; neoliberalism; mechanism design; markets

Government contracting is ubiquitous and unavoidable. But there are interesting questions about the when and the why. In some cases, contracting government services to private-sector firms, for-profit or not, is straightforwardly impermissible, even if it results in efficiency gains. For example, arguments against privatization from broadly Kantian first principles concentrate on the wrongs of contracting out core coercive functions of the state such as defence and criminal justice to mercenaries and private prisons (Pattison 2010; Dorfman and Harel 2013, 2016; Farrell 2019; Harel 2019; Cordelli 2020).

Concentrating on the worst cases, as the Kantians do, obscures the variety of both reasons for contracting and goods and services that are contracted. For example, governments often contract for the same reasons that drive firms to buy goods and services rather than making them in-house. Sometimes, those reasons justify contracting, and those decisions are examples of what Heath (2023) rightly calls 'anodyne privatization'. For example, when public employees – everyone from military service members traveling on orders to public university professors going to conferences – travel, they generally fly commercial, and the government pays for it. This sort of privatization of public-service travel seems obviously superior to the silly alternative of the North Carolina Air National Guard's C-17 squadron ferrying UNC's faculty around. Unless you're a proponent of thoroughgoing state socialism, you have to admit that the government buying services from private companies

instead of performing them in-house isn't intrinsically objectionable. And so far as I can tell, nobody has much of a problem with this. Heath rightly highlights that, by contrast to its more objectionable forms, this sort of privatization is so normatively unobjectionable that nobody really notices it.

The Kantian first-principles argument against privatization and Heath's 'anodyne' privatization are both interesting models of privatization, but obviously neither is comprehensive. I won't pretend to be either. But I'm going to argue for a third normative category, motivated and justified by similar reasons to the ones that justify paradigmatic liberal institutional forms such as property rights protections and the welfare state. Part of the point of those is to reduce our vulnerability to one another (Shklar 1989, 2023). Indeed, the worry that privatization exacerbates that vulnerability by turning the business of coercion over to unaccountable private entities often motivates scepticism about privatization (Cordelli 2020).

Technological innovations such as vaccines and carbon capture and sequestration technologies can also help mitigate vulnerability. Unfortunately, they can fall victim to a specific kind of market failure called the 'hold-up problem'. Governments can remedy this using a specific contracting mechanism called an 'Advance Market Commitment', which is basically a promise to buy some amount of a specified good at a pre-specified price, from whomever can provide it. AMCs use the government's capacities as a buyer to create a market for socially valuable goods where none existed before. They're a non-standard policy tool that allows governments to help mitigate citizens' vulnerability, specifically by buying from the private sector. I aim to show that they're neither objectionable from Kantian first principles, nor entirely anodyne, because they help bring socially useful technologies to market.

1. Make or Buy?

One underappreciated dimension of privatization, at least among political philosophers, is that the government often has to make decisions whether to produce some good or perform some service in-house or buy it from an outside provider. Heath rightly notes that most of the reasons that motivate those decisions are similar to the ones that motivate private companies. I want to lay out a detailed case, just to get clear on the kinds of considerations Heath highlights. Afterwards, we'll see how AMCs can have more complicated and interesting normative features than paradigmatic cases of governments choosing to buy instead of make.

It's important to note at the outset that 'make' and 'buy' are ideal types. The Kantian critics of privatization tend to concentrate on a small subset of these decisions that involve replacing the government with a private entity. But many interesting cases involve a mix of the two, and it makes more sense to think of particular cases as points on a spectrum. I want to start by laying out the traditional make-or-buy decision to contrast it, further on, with the more normatively rich approach to contracting AMCs help us see.

NASA made make-or-buy decisions, for different reasons, between the 1950s and the present. The early Mercury, Gemini, and Apollo programmes, the contemporary Space Launch System, and the 2006–2013 Commercial Orbital Transportation

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

Services/Commercial Resupply Services programmes all represent different points along the make-or-buy spectrum. Much of the work for the first was done entirely within NASA's Marshall Spaceflight Center. The second is largely funded by traditional 'cost-plus' or 'push' funding contracts to contractors such as Boeing and Northrop Grumman that pay for inputs to the development process, combined with tight control over design features by the contracting government agency. The last is an AMC for rocket launch services.

The point of the comparison isn't to give a definitive judgement that one is better than the other: the Apollo and COTS/CRS programmes were both successful, whereas the Space Launch System has more or less succeeded only in lighting tax dollars on fire (NASA Office of Inspector General 2021a: 19–25; 2021b, 2022, 2023c). It's just to highlight in some detail the different considerations that can motivate and (more importantly for our purposes) justify a particular make-or-buy decision.

The early American space programme included the development of launch vehicles that led to landing an astronaut on the moon. Many, including the Saturn V rocket, were designed in-house by engineers at Marshall Space Flight Center in Alabama, employed by the United States government (Dunar and Waring 1999: 39–78; Kranz 2000: 19). As NASA's official history of the Center puts it, '[r]ather than being primarily supervisors of contractors, Center personnel were hands-on designers, testers, manufacturers, and operators' (Dunar and Waring 1999: 39). This was, in large part, due to the lack of any developed rocketry industry in the United States (or anywhere else). The initial launch platforms were barely modified Redstone intercontinental ballistic missiles, whose limitations as a launch platform quickly led to the development of better rockets. Because nobody had ever done this before, NASA wanted near-absolute control over each stage of the design process. Designing, building and testing new technologies under one roof allowed close coordination between different components of the design process. While contractors certainly played a role in building those components, the design expertise was mostly housed within Marshall (Dunar and Waring 1999: 40–43).

Optimal programme design requires navigating some trade-offs. The 'privatization' of this particular function of the state is a constant negotiation between various competing factors given to us by the structure of the state and the structure of industry. In the 1950s and 1960s it made sense to develop much of the technology in-house, if for no other reason than the lack of expertise in the private sector. However, the contemporary landscape is complicated by the fact that the government pays less than private companies, who can afford to hire top-level engineering talent, which makes contracting attractive, if not unavoidable. But the management of contracting relations requires decisions about the division of labour between NASA and its contractors. Those decisions are basically motivated by efficiency. Managers want, in principle, to get NASA the biggest bang for the taxpayer's buck, subject to the constraints on NASA's own capacities.

The traditional method of procurement from contractors uses 'push' funding: paying for inputs rather than a final product. Governments often use push funding because it allows them to control detailed aspects of development and manufacturing. The Space Launch System programme, designed to support the Artemis programme to return astronauts to the moon, is largely being developed by a conventional push-funded contract between NASA and a consortium made of

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

Northrop-Grumman, Boeing and subcontractors (NASA Office of Inspector General 2023b: 5–6).

Unfortunately for the causes of space exploration and responsible stewardship of Americans' tax dollars, but fortunately for the cause of a paper arguing that AMCs are useful and interesting, it has exhibited many of the pathologies of traditional push-funded contracting. For example, NASA's Office of Inspector General reports projected SLS costs of about 60% over initial projections between 2014 and 2020, or up to $2.5 billion per launch (NASA Office of Inspector General 2020: 15; 2023a: 1). Part of the problem is that the structure of the SLS' contracts is opaque, which makes tracking the costs of individual deliverable items difficult. Much of this is due to the internal organization of NASA's contracting mechanisms. For example, the offices responsible for particular components of the broader Artemis programme often don't communicate with one another about contractors' performance. This failure means that NASA is often unable to anticipate potential cost overruns or delays. What's more, NASA itself can't generate a list of every supplier it uses. In other words, the offices at NASA responsible for managing contracts lack the capacity to realize the advantages of cost-plus contracting (NASA Office of Inspector General 2020: 18; 2023b: 12–15; Government Accountability Office 2023).

The Inspector General basically gives two sorts of recommendations to remedy this sorry state of affairs. First, it recommends reorganization of the contracting and procurement bureaucracy to allow better flow of information between programmes. This is the 'make' side of the 'make-or-buy' decision. It requires building some infrastructure within NASA to keep track of suppliers and their performance, allowing programme managers to proactively address potential supply-chain and other procurement issues. It also might require collaborating with other agencies inside the US government who do a better job keeping track of their contractors' supply chains.

Second, it recommends, in effect, further privatization of the SLS development effort by transitioning to fixed-price contracts and greater use of commercial services. Writing a fixed-price contract and committing to particular performance benchmarks in advance is a way the government cedes control over the development process to the contractor. It takes processes that happen *inside* the government on the traditional contracting model and effectively black-boxes them inside the contractor. It thus sacrifices visibility inside contractors' operations. But it allows the government to track its costs a bit more precisely than the status quo because it spells out what the government will pay for specific deliverables (NASA Office of Inspector General 2023a: 20). This is what Heath (2023: 45) calls a 'Type 2' privatization, in which 'goods that had once been made internally are now bought from external sources'. In effect, the 'make-or-buy' decision trades off control for costs, subject to the various capacities of the NASA procurement bureaucracy.

The extreme case of this sort of contracting is NASA's Commercial Orbital Transportation Services (COTS) programme, which was effectively an advance market commitment for rocket launches into low-earth orbit. We'll concentrate on the fact that an AMC is a kind of fixed-price contract just to establish that it is a kind of privatization, and then worry about the normatively interesting aspects of it later.

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

Between 2006 and 2013 NASA solicited bids for private companies to develop rocket technologies that could be used to resupply the International Space Station and carry payload into orbit both for NASA and for other customers. What's interesting about the COTS programme was the structure of the contracts. NASA selected multiple 'winning' bids from industry, and agreed to pay fees agreed upon in advance for the completion of particular milestones. By committing up front to a set of milestones, NASA ceded control over the design and development process to their commercial partners (NASA 2014: 106–108). And the follow-on Commercial Resupply Services (CRS) contracts similarly paid an agreed-upon price for discrete services: orbital launches carrying a particular payload on privately owned and operated rockets. For example, SpaceX was paid $1.6 billion for 12 cargo launches on rockets *that it owned*. By contrast, previous NASA projects to send people and things into space used equipment such as the Space Shuttle and the Saturn V rocket owned by the taxpayer and operated by NASA. One major payoff is cost. NASA estimates the entire development cost for SpaceX's Falcon 9 rocket at about $400 million, or about a tenth of the counterfactual scenario where that development was funded conventionally. Another is that COTS was explicitly designed to create a market for, and thus spur the development of, commercial, private-sector space launch capabilities (NASA 2014: 93–98; Congressional Research Service 2016).

So we've seen a taxonomy of ways government programmes can be more or less privatized. Rockets can be designed and produced entirely within the confines of the Marshall Space Flight Center; bought from a private contractor who nevertheless responds to detailed design guidance from the government; or bought on a fixed-price contract from a vendor who meets certain design specifications, but other than that cedes virtually no design control to the government. This helps establish that fixed-price contracts, of which AMCs are a variety, count as privatization relative to the contrast classes of the Mercury and Space Launch System programmes because they allow the contractor authority over decisions that other contracting paradigms keep within the government. They're not privatization in the pejorative, Kantian sense, in that they don't seem to illegitimately transfer power away from the public sector. The motivation for this sort of privatization is largely to make the government's own activities more efficient. But that's obviously not the only reason governments decide to contract or not. They might, for example, accept efficiency losses in order to avoid compromising national security.[1] In the next section, I'll argue that we should understand the motivation for AMCs on *liberal* grounds, rather than narrowly efficiency-promoting ones. The problems they solve are normatively similar to problems solved by things such as property rights and the welfare state.

2. Not Just Economizing

AMCs use the government's ability to buy stuff to remedy market failures. That's why they're normatively interesting. I'll show in this section that the market failures they're designed to solve can undermine important liberal values. Liberal institutional forms, in general, aim, at least in part, to reduce our vulnerability to one another. Whatever internecine disagreements liberals might have, they tend to

[1]Thanks to an anonymous referee for a helpful discussion of this.

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

agree on this. Liberal institutions aim at 'freedom from the abuse of power and intimidation of the defenseless' (Shklar 1989: 27). That core normative assumption entails recognizing that public institutions ought to aim at mitigating our vulnerabilities, especially to one another.

Apart from the paradigmatic forms of cruelty, many social ills exhibit a phenomenon called 'tight coupling'. This is a characteristic of the interactions between component parts of a system. Tightly coupled systems are vulnerable to cascades of failure stemming from the failure of individual parts (Perrow 1999: 89–100).[2] So if you're a subsistence farmer, a bad harvest, an attack of some disease (either on you or your crops), some rats getting into your grain cellar because you've neglected to seal it properly, or any other failure of some part of your life brings the risk of utter catastrophe. Similarly, if you're a working-class college student, any single failure – say, if you have to care for an ailing parent – can completely derail your ability to finish college and limit your future economic prospects in a way that a richer student might not face (Cook 2019). In other words, tight coupling increases vulnerability to cruelty. Loosely coupled systems, on the other hand, exhibit slack between their component parts.

One way liberal politics seeks to relieve vulnerability is by introducing that slack. Perhaps the paradigmatic example of this is the modern welfare state's role as social insurance. Insurance works through a principle of 'risk pooling'. People pay an amount into a common fund on the understanding that they'll receive a comparatively large payout in case of some reverse. This effectively decreases the variance of whatever outcomes might result from their various pursuits. This just is the value proposition of the welfare state. While most contemporary defences of the welfare state see it as a way to remedy inequality, it's hard to make sense of things like unemployment insurance except as *insurance* – a way to mitigate the risks of becoming unemployed. In combination with a hot labour market, unemployment insurance reduces vulnerability both to domination by one's boss (because it makes quitting less costly) and to the prospect of poverty and deprivation (by making it easier to live without a source of income) (Heath 2006: 322–324; 2011). The various Great Society programmes, as well as Social Security, similarly work by letting the elderly and indigent afford medical care and living expenses – hence Krugman's (2011) quip that the US government is 'best thought of as a giant insurance company with an army'. The welfare state, like all insurance, helps decouple the shock of losing a job from other bad consequences like losing your house or the ability to afford food or medical care. It's important to emphasize that the sources of vulnerability that the welfare state guards against can be natural, social or some combination of the two. A functional welfare state is useful to you whether you lose your job or lose your house to a hurricane.

Liberal rights can be justified on similar grounds. They are, for Shklar, 'an indispensable and excellent way of limiting the long arm of government and of dividing social power, as well as of securing the independence of individuals' (1989: 31). This

[2]The philosopher Iris Marion Young has named a similar phenomenon, involving the interactions of many individually blameless components, 'structural injustice' (Young 2011), and while I understand the presumption against introducing jargon for its own sake, I prefer the terminology of 'tight coupling' because it's a bit more precise.

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

theme persists across differences in background moral commitments within liberal thought. There are plenty of disagreements within liberalism about the moral basis for rights protections. I just want to establish that the theme of mitigating vulnerability persists across those disagreements.

So, for example, Hardin (1988: 78) gives an explicitly utilitarian basis for rights protections. He argues that they solve information problems. This is important because without rights, control over important parts of individuals' lives might default to those who don't know what's best for a particular individual. Better, then, to have individual spheres of sovereignty delineated by rights protections, because even in cases where, intuitively, violating rights increases utility, systematizing and institutionalizing a system to tell us when those violations would increase utility would be *extremely* difficult at best, and at worst intrusive in a way that would undermine the utilitarian basis of the project in the first place. In other words, Hardin's utilitarian defence of rights protections is about mitigating the often unintended and often negative consequences of people acting in situations where they lack information. Even if you don't go in for utilitarianism, this seems like a worthwhile way to think about liberal rights.

While Hardin was most concerned to reconcile the protection of rights with utilitarian moral theory, Jerry Gaus (2011: 372–386) defended them on the grounds that they help us economize on the costs of coming to an agreement about some social rule. They keep disagreement from spilling over into oppression. One of the major themes of political philosophy over the past half century has been the importance of public justification for legitimacy. The moral considerations you think ought to bear on some area of public concern are worth no more than anyone else's. But it'd be a *very* odd and very oppressive society in which everything is an area of public concern, subject to public justification. Jurisdictional rights, then, are a respite from hegemonic public scrutiny.

Regardless of your underlying theory of the moral basis of liberalism, it makes sense to think of the protection of rights as a way to mitigate the vulnerabilities that come from tight coupling in social and political life, whether that's due to information costs or the burdens of public justification. The innovations that AMCs are designed to deliver do something similar. Besides rockets, AMCs have been used to buy vaccines for pneumococcus (Berndt *et al.* 2005, 2007; Kremer *et al.* 2020) and (most notably) COVID-19 (Slaoui and Hepburn 2020; Nocera and McLean 2023), as well as carbon-capture and sequestration technologies (Athey *et al.* 2021; Department of Energy 2023). Each of these innovations is valuable because it helps mitigate our vulnerability to one another, by helping loosen tightly coupled social structures. The problems these innovations solve all give rise to the sorts of vulnerability liberal institutional forms seek to mitigate.

For example, climate change has all kinds of terrible compounding effects. High heat lowers cognitive performance (Park *et al.* 2020), raises crime rates (Heilmann *et al.* 2019), increases the risk of natural disasters, may have terrible effects on agricultural productivity, and lowers economic growth (Kahn *et al.* 2021), all of which are bad from the perspective of liberal politics. And those effects interact with the various other sources of injustice. If you're already poor, they're only going to immiserate you further. Poorer areas of American cities, for example, tend to be hotter. And those areas obviously tend to be less well resourced, have worse schools,

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

and be more violent at baseline. All this is to say that climate change exacerbates a number of existing inequalities and injustices. They are parts of a tightly coupled system analogous to the consequences of poverty that the welfare state is supposed to help remedy.

Carbon capture and sequestration technology allows us to decouple, at the margin, dirty sources of energy such as oil and natural gas from the harms caused by climate change. They also allow us to eliminate historical carbon dioxide emissions, undoing the damage of the Industrial Revolution. What makes them normatively interesting for our purposes is that they make us less vulnerable to the problems associated with climate change. They protect us from some of the worst consequences of our current state of energy technology. And they thus reduce some of the vulnerabilities, both natural and social, that climate change exacerbates. And because CCS isn't excludable, those benefits are more egalitarian. Poor people can enjoy them just as much as rich people.

Infectious diseases have similar features. Even when they're not debilitating, they raise the relative costs of the various activities that lead to a flourishing life, and thus interact with and exacerbate various sources of disadvantage that we care about. The COVID-19 pandemic, for example, disproportionately affected lower-income Americans because working remotely is easier in higher-income professional jobs held by more educated workers (Piacentini *et al.* 2022: 12–13). Pandemic-related learning losses were much worse in minority and low-income districts than they were in higher-income districts (Goldhaber *et al.* 2023). That is, they exacerbated existing inequalities in math and reading proficiency. In the developing world, treatable diseases can, just like crop failure, lead to widespread death. If your life is already quite precarious, the incidence of disease, even absent a global pandemic, is likely to make things worse (Pinto 2023).

Immunizations are important because they help make poor people's lives less precarious. Immunization campaigns are a central part of public health, especially in developing countries, because mitigating disease incidence is so helpful for all the other goals a comprehensive development programme might have. So, the rollout of the COVID-19 vaccines in the USA have allowed for school and socializing to continue relatively uninterrupted, and while remote work has become more popular, it at least allows for the resumption of working in offices (Gupta *et al.* 2022). It also makes jobs where workers have to show up in person *much* less dangerous. Similarly, in the developing world, widespread vaccination can allow people to work and send their children to school in the face of disease outbreaks that might otherwise have prevented it. There is a worthwhile analogy with programmes like Medicaid in the USA. Medicaid is designed to make it easier for poor Americans to access healthcare, and thereby mitigate the effects of poverty on the vulnerabilities caused by illness. Developing new vaccines helps serve the same goal by mitigating everyone's vulnerability to illness.

3. How AMCs Solve the Hold-Up Problem

Lots of technological innovations are valuable for the same reasons as paradigmatic liberal institutional forms. Unfortunately, many innovations face a particular kind

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

of market failure called the *hold-up problem*. It's worth going into some detail about why this problem arises, why it undermines important liberal values, and how AMCs help solve it. This helps show that public sector procurement from the private sector serves those core liberal goals of mitigating vulnerability.

Obviously the private sector does plenty of research and development on its own. But sometimes socially useful innovations are never widely adopted. This form of market failure happens when a firm's decision to invest in new technology confronts an uncertain demand signal. Imagine, for example, that you're the CEO of a pharmaceutical or biotechnology company deciding whether to invest in higher production capacity for a useful new vaccine, or a venture capitalist deciding whether to invest in a start-up developing a new carbon capture and sequestration technology. Part of what you want to know is whether you'll be able to sell the product you're considering developing. If you can't, then your investment will be lost, your shareholders will get annoyed, and (most importantly) you'll have failed to invest your money somewhere more useful.

How will you know whether your investment is likely to pan out? Roughly, you have to be able to produce the new good at a unit price *below* what buyers are willing to pay. However, below a certain unit price (called your *reservation price*), you're better off just making a different investment. So you have to hope that the buyer's private valuation is above your reservation price. The catch is that a buyer who knows that you've invested in a new technology can bargain down the unit price to appropriate the entire surplus of the transaction (Kremer *et al.* 2022: 4790; Vanderschraaf 2023). In equilibrium, without a buyer who can commit to paying more than your reservation price for the new technology, you're likely to try to protect yourself against the risk of your investment failing to break even by investing less than the socially optimal amount into the new technology (Bueno de Mesquita 2016: 180–183). This undermines our ability to serve the liberal goal of mitigating vulnerabilities.

Hold-up problems are ubiquitous in contracting, both in the public and private sector. But they're especially thorny when the goods in question have external benefits and display learning curves. While there are thorny normative questions about which goods ought to be contracted – in other words, what's the in principle difference between vaccines and mercenaries? – I think there are good reasons to think that the combination of positive externalities and learning curves is a good first cut, because they seem to be responsible for the specific market failures that AMCs are designed to solve.

3.1. Positive externalities

An externality is a cost or benefit produced by a good's consumption that accrues to someone other than the consumer. Carbon capture and sequestration technology is an example of a good with social benefits that can't be internalized by a buyer. Nobody particularly wants to buy captured carbon dioxide. Mostly you just try to find a way to squirrel it away underground. On the other hand, if someone else buys a bunch of sequestered carbon dioxide, everyone else benefits from the resulting climate change mitigation. Obviously the benefits of a cooler planet aren't excludable. This incentivizes everyone to free-ride on everyone else's presumptive

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

purchases of captured carbon dioxide. Because the costs of capturing and storing carbon dioxide are so concentrated and their benefits diffuse and non-excludable, the amount that buyers might be willing to pay is above presumptive sellers' reservation price, which 'holds up' investment to a level well below the social optimum.

Space launch technology is similar. During the early days of the space programme, nobody beyond NASA and the military had any use for launching a payload into low-earth orbit. Things have changed in the intervening three-quarters of a century. Today, access to low-earth orbit underpins all sorts of vital technologies. The Global Positioning System we all use to find our way around requires satellites. New internet access technologies use satellites. NASA has a system called FIRMS that uses satellites to monitor outbreaks of fires on the earth's surface. FIRMS and similar systems have also been used by news organizations to track the war in Ukraine, because they detect the muzzle flashes of artillery, the impacts of airstrikes, and the like (Gonzales 2022). Cheaper launch costs have major knock-on benefits.

Vaccines – the original use case for AMCs (Cernuschi *et al.* 2011; Kremer and Glennerster 2004; Dattani *et al.* 2023) – also exhibit positive externalities. The point of being vaccinated against an infectious disease is, in part, to protect yourself from being infected by someone else, and, in part, to protect others from being infected by you. The benefits of being vaccinated aren't entirely internalized, and they aren't excludable – you can't choose which people get protected by your decision to get vaccinated. So, for example, COVID-19 vaccines are estimated to have a social value in the neighbourhood of between hundreds and thousands of dollars per course (Castillo *et al.* 2021: 1107; Glennerster *et al.* 2024). This is orders of magnitude higher than their commercial value. Similarly, the early deployment of next-generation pneumococcus vaccines through the GAVI AMC has been estimated to have saved about 700,000 lives through 2020, relative to a baseline without the AMC (Kremer *et al.* 2020: 269). No group of individuals will be able to internalize the full social value of a vaccine, so economic theory predicts that markets will undersupply them.

3.2. Learning curves

Some goods also exhibit 'learning curves' – that is, their unit production costs decline as a function of units produced (Nemet 2019; Malhotra and Schmidt 2020). If you have to build a bunch of something in order to be able to figure out how to make them cheaply enough to sell them at a price where someone will buy them, then you'd better find a way to pay to make the first bunch of units.

The situation isn't totally hopeless. It's not as though technology that's too expensive to be practical is always *impossible* to sell. In the *Wealth of Nations* (1981 [1776]: III.iv.10–11), Adam Smith, for example, credits the vanity of the idle rich and the resulting demand for tchotchkes[3] with spurring the development of commercial society. Thorstein Veblen (1899) similarly argues that status anxiety leads to competitive consumption by the wealthy. And at least in some cases, the

[3]Not his word, obviously.

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

private benefits of having the newest toy will spur early adopters to buy enough of a product to fund the learning-curve process. For the goods at issue here, though, the fact that the benefits of consumption are hard to internalize makes it doubly hard to find any kind of demand at the top of the learning curve,[4] before unit production costs have fallen significantly. The units that have to be produced in order to realize efficiency gains may well cost more to produce than the private valuation of consumers. This means that potentially useful technologies might never be deployed en masse.

3.3. Punchlines

So there are good reasons to believe that both positive externalities and learning curves lead to hold-up problems. Investing in a new technology is expensive and risky. In the absence of a credible signal of demand, these investments get riskier. Private capital becomes harder to access. Firms at the bleeding edge of socially useful technological innovation will likely recognize the risks of research and production capacity investments in the presence of weak demand signals. They will thus underinvest, leading to less than socially optimal amounts of innovation or production. And this is one explanation for why basic microeconomic theory predicts that goods with positive externalities – like innovation – will be undersupplied by a market.

Vaccines are the best example of the hold-up phenomenon. They're the most obvious case of socially necessary goods with massive external benefits. They are extremely expensive to produce, so there's a question about how to provide them to poor countries with ill-funded public health apparatuses. They also operate on a learning curve – manufacturing has very high fixed costs, but making more doses unlocks efficiencies that allow lower unit costs for successive doses. Without some guarantee of demand, installing enough vaccine capacity to drive manufacturing costs down is quite risky. Those risks are also socialized. If it can't sell enough doses to recoup its investment, a company might face financial problems. But more importantly, the capital might have been expended on other things. Wasted pharmaceutical investment costs lives. So just like an inadequate welfare state, suboptimal investment in pharmaceuticals makes it harder to mitigate important vulnerabilities.

In the model of the hold-up problem I sketched above, buyers can't commit to paying a price above the seller's reservation price. But in some cases, real-world buyers can write contracts that commit them to buying a certain number of units at a higher price. This is the basic logic behind an AMC. It's the 'commitment' in 'Advance Market Commitment'. Effectively, AMCs provide a per-unit subsidy to producers, tied to the proportion of overall need each producer provides (Kremer 2000a: 45–46; Kremer 2000b; Levine *et al.* 2005; Kremer and Williams 2010; Kremer *et al.* 2020: 270). So if a funder commits to a price of $5/dose for a total of one billion doses, a supplier who can commit to a hundred million doses would be entitled to half a billion dollars in purchases. Unlike a prize, AMCs aren't create winner-take-

[4]Bottom of the learning curve? It's hard to know what the right spatial metaphor is. What I mean is at the *beginning* of the learning curve, before unit production costs have been drastically reduced.

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

all tournaments (Frank and Cook 2013; Howes 2022). Instead, they create a market for a good or service that supports multiple producers. Operation Warp Speed, for example, funded a total of eight vaccine candidates (Slaoui and Hepburn 2020: 1702–1703). COTS awarded contracts to two different companies, SpaceX and Orbital Sciences Corporation (NASA 2014: 50, 64). And the climate philanthropy Frontier (2022) explicitly buys stored carbon from anyone who meets certain technical criteria. Note that while there is a worry that an AMC for some good might not give rise to a competitive market, at the margin it's less vulnerable to worries about monopoly power than other kinds of innovation policy. For example, pharmaceutical development is incentivized with *patents*, which explicitly grant a monopoly for a new drug. AMCs improve on this both by incentivizing multiple suppliers, and by paying for those suppliers to supply new drugs at lower prices (Kremer and Glennerster 2004).

AMCs solve the hold-up problem because they allow producers to anticipate being able to sell more units for more revenue, which incentivizes a higher investment level. Under Operation Warp Speed, the federal government agreed to buy one hundred million doses of the Pfizer-BioNTech COVID-19 vaccine for just under $2 billion, with an option to buy another half-billion doses, conditional upon FDA approval. The agreement allowed Pfizer, which was confident in its eventual approval, to immediately begin to scale up production in order to meet demand (United States Department of Health and Human Services 2020). This, obviously, was important given the time pressures on vaccine deployment. In effect, by paying to scale up manufacturing in parallel with clinical trials, the federal government assumed the risk of one, some, or all of the candidate vaccines failing their clinical trials (Slaoui and Hepburn 2020: 1702). By guaranteeing purchase of hundreds of millions of doses of a vaccine upon FDA emergency use authorization, the federal government defrayed the risk of the investment needed to scale up production facilities during the dark days of summer and autumn 2020.

To be clear, Operation Warp Speed involved more than just an advance market commitment. It also included a significant amount of push funding to various candidate vaccine manufacturers, including Moderna; a significant intervention into the logistics of vaccine production by the Department of Defense; a slightly different clinical trial structure; and other interventions (Pfizer accepted no funding from the federal government other than the purchase of vaccines). But the advance market commitment helped defray risk from companies' decisions to pour their own resources into making vaccines. If the government had not created the entire demand side of the market for COVID vaccines, the Pfizer and Moderna vaccines wouldn't have come to market as quickly as they did (Nocera and McLean 2023).

4. Contracting for Justice

The major value proposition of AMCs in a liberal society is that they remedy market failures that exacerbate vulnerability. This requires showing, first, that *contracting* is what fixes those market failures, and second, that AMCs don't run afoul of common criticisms of contracting. This isn't the same as arguing that AMCs are the optimal

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

or ideal policy solution, relative to some counterfactual ideal political reality. I just want to show that a particular form of contracting by governments can advance liberal values, without taking the stronger position that it's *required* by any conceivable form of liberalism.

One reason to think contracting is a useful mechanism for remedying market failure is that the government's ability to credibly commit to purchasing goods and services has some interesting knock-on effects. For example, the COTS contracts incorporated an agreement to purchase services, but also required potential suppliers to acquire significant investment on their own specifically to ensure that they weren't simply wasting NASA's resources (NASA 2014: 106). The contract design helped 'crowd-in' private investment by providing a guaranteed return on investment, conditional upon the supplier's ability to successfully supply the required goods. In other words, it makes investment *much* less risky. The idea was to create a market with multiple suppliers for socially valuable goods who might, in at least some cases, go on to commercial success. Whether or not a market for some product that includes non-government entities on the demand side actually develops depends on the nature of the products or service in question. If enough of the benefits of consuming a good are sufficiently internalizable, then such a market might develop. The complementarity between public and private financing helps do that. A good example of this is the cost of low-earth orbit launch services, which has fallen by an order of magnitude in the past 15 years, thanks in large part due to the cycle of development and refinement kicked off by SpaceX's participation in the COTS/CRS programmes (Berger 2021; Roberts 2022).

The core normative case for AMCs proceeds by analogy with other kinds of policy that remedy market failures. Much of the welfare state in developed countries takes the form of an insurance product supplied by the government. In the USA, Medicaid is a government-run health insurance programme; similarly, every state has an unemployment insurance programme. The point of these programmes is to serve as an 'insurer of last resort', to remedy the adverse selection problems that plague private health insurance markets (Heath 2011: 34–37; Einav and Finkelstein 2023). So the welfare state makes us less vulnerable to each other specifically by solving a particular kind of market failure.

AMCs have a similar justification. The problem of missing markets for socially useful goods and services is a kind of market failure. If you accept Pigovian subsidies in principle, then there's no good reason not to accept AMCs. Consider the alternative of the government subsidizing consumers instead of buying a good directly. By hypothesis, all the interesting applications of AMCs involve goods that exhibit positive externalities, so the sum total that consumers will be willing to pay will be lower than the social value of the good. What's more, the commitment required to resolve the hold-up problem creates one *hell* of a collective action problem among millions of consumers. This entire collective action problem can just be resolved by having the government do all the buying, instead of subsidizing demand.

AMCs also usefully don't fall afoul of any of the traditional criticisms of government contracting with the private sector. This includes the criticism from basic Kantian first principles that 'privately achieved justice is not justice' that turns on a tight conceptual link between justice (or whatever important value) and

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

specifically *public* governance/provision/adjudication/what have you. So, for example, Cordelli argues that only a particular sort of political order – the jargon here is 'civil condition' – can square the idea that when we make claims of justice (that is, about the contents of our and others' entitlements and obligations) we change others' entitlements and obligations too with a commitment to respecting 'the equal normative authority of all and … individuals' rational independence' (Cordelli 2020: 47). Respecting those two baseline normative commitments requires making sure that nobody is subject to the unilateral will of another, and Cordelli's central claim is that private organizations cannot do that. Similarly, Dorfman and Harel both argue that privatization is bad because it replaces some functions that should – and this is a moral 'should' – be performed by public officials with analogues performed by private actors. Privatization 'severs the link between decision-making processes and citizens, eroding the prospect of meaningful political engagement and civic shared responsibility' (Harel 2019: 53). This undermines the democratic authorization of those actions, and thereby undermines their legitimacy.

Again, this is a conceptual rather than empirical point. It's not that public organizations are uniquely capable of solving certain kinds of problems that, left unsolved, exacerbate vulnerability to unilateral wills. The point is that what it is for a political state of affairs to properly count as *not* subjecting people to the unilateral will of others is for it to have a set of properties that privatized goods and services *don't* have, in virtue of what it is to be privatized (Cordelli 2020: 17).

This all takes place at an *extremely* high level of abstraction – as it has to, if the claims are to have the reach that's claimed of them. But conceptual critiques all argue that privatization constitutively involves the hollowing out of the state's apparatus. It involves the *replacement* of some function performed in-house by one performed by a contracted private-sector entity. Privatization, as an aggregate phenomenon, is just the entire set of such replacements. And making the inference from examples of privatization to aggregate characteristics requires either assuming or arguing that all, or at least most, examples of privatization (or, more broadly, contracting by the government) contribute to undermining collective self-government.

AMCs don't run afoul of this sort of objection. While they count as a form of privatization, in that they involve the government buying things from the private sector instead of building them in-house, it's hard to see how they subject anyone to a unilateral will. In large part this is because the market failures they help remedy are ones that make us more vulnerable to others' wills. Similarly, it's hard to see how they constitute the *replacement* of a legitimate government function by an unaccountable private entity. If AMCs run afoul of this sort of criticism, then so does *any* government procurement from the private sector. Anyone committed to it has to embrace thoroughgoing state socialism, at which point they've got bigger problems. The conclusion to draw, then, is that advance market commitments are a form of contracting that is normatively distinct from the ones that motivate most extant philosophical concerns about privatization.

Acknowledgements. I'm grateful to Abe Singer, Ryan Muldoon, Jake Monaghan, Chris Snyder, Joe Heath, Russ Muirhead, Max Tabarrok and Kris Gulati, as well as two anonymous referees, for helpful discussion.

Financial support. The author declares no financial support or conflicts of interest.

References

Athey S., R. Glennerster, N. Ransohoff and C.M. Snyder 2021. Advance market commitments worked for vaccines. they could work for carbon removal, too. *Politico*. www.politico.com/news/agenda/2021/12/22/carbon-removal-advance-market-commitments-525988.

Berger E. 2021. *Liftoff: Elon Musk and the Desperate Early Days that Launched SpaceX*. New York: HarperCollins.

Berndt E.R., R. Glennerster, M.R. Kremer et al. 2007. Advance market commitments for vaccines against neglected diseases: estimating costs and effectiveness. *Health Economics* **16**, 491–511.

Berndt E.R. and J.A. Hurvitz 2005. Vaccine advance-purchase agreements for low-income countries: practical issues. *Health Affairs* **24**, 653–665.

Bueno de Mesquita E. 2016. *Political Economy for Public Policy*. Princeton: Princeton University Press.

Castillo J., A. Ahuja, S. Athey et al. 2021. Market design to accelerate COVID-19 vaccine supply. *Science* **371**, 1107–1109.

Cernuschi T., E. Furrer, N. Schwalbe, A. Jones, E.R. Berndt and S. McAdams 2011. Advance market commitment for pneumococcal vaccines: putting theory into practice. *Bulletin of the World Health Organization* **89**, 913–918.

Congressional Research Service 2016. Changes in the US Commercial Space Industry. https://crsreports.congress.gov/product/pdf/IF/IF10415.

Cook G. 2019. The economist who would fix the American dream. *The Atlantic*. www.theatlantic.com/magazine/archive/2019/08/raj-chettys-american-dream/592804/.

Cordelli C. 2020. *The Privatized State*. Princeton: Princeton University Press.

Dattani S., R. Glennerster and S. Haria. 2023. *Why we didn't get a malaria vaccine sooner?* Works in Progress. https://worksinprogress.co/issue/why-we-didnt-get-a-malaria-vaccine-sooner/.

Dorfman A. and A. Harel 2013. The case against privatization. *Philosophy and Public Affairs* **41**, 67–102.

Dorfman A. and A. Harel 2016. Against privatisation as such. *Oxford Journal of Legal Studies* **36**, 400–427.

Dunar A.J. and S.P. Waring 1999. *Power to Explore: A History of Marshall Space Flight Center 1960–1990*. NASA Office of Policy and Plans. www.nasa.gov/wp-content/uploads/2023/02/SP-4313.pdf.

Einav L. and A. Finkelstein 2023. *We've Got You Covered: Rebooting American Health Care*. New York: Portfolio/Penguin.

Farrell H. 2019. Privatization as state transformation. *NOMOS* **60**, 171–199.

Frank R.H. and P.J. Cook 2013. Winner-take-all markets. *Studies in Microeconomics* **1**, 131–154.

Frontier 2022. Stripe, Alphabet, Shopify, Meta, and McKinsey launch advance market commitment to buy $1B of carbon removal by 2030. https://frontierclimate.com/writing/launch.

Gaus G. 2011. *The Order of Public Reason*. New York: Cambridge University Press.

Glennerster R., T. Kelly, C.T. McMahon and C.M. Snyder 2024. Quantifying the social value of a universal COVID-19 vaccine and incentivizing its development. NBER Working Paper 32059. www.nber.org/papers/w32059> doi:10.3386/w32059.

Goldhaber D., T.J. Kane, A. McEachin, E. Morton, T. Patterson and D.O. Staiger 2023. The educational consequences of remote and hybrid instruction during the pandemic. *American Economic Review: Insights* **5**, 377–392.

Gonzales C. 2022. Scorched Earth: using NASA fire data to monitor war zones. *Bellingcat*. www.bellingcat.com/resources/2022/10/04/scorched-earth-using-nasa-fire-data-to-monitor-war-zones/.

Government Accountability Office 2023. Space launch system: cost transparency needed to monitor program affordability. www.gao.gov/assets/gao-23-105609.pdf.

Gupta A., V. Mittal and S. van Nieuwerburgh 2022. Work from home and the office real estate apocalypse. NBER Working Paper 30526. doi: 10.3386/w30526.

Hardin R. 1988. *Morality Within the Limits of Reason*. Chicago: University of Chicago Press.

Harel A. 2019. Why privatization matters. *NOMOS* **60**, 52–78.

Heath J. 2006. The benefits of cooperation. *Philosophy and Public Affairs* **34**, 313–351.

Heath J. 2011. Three normative models of the welfare state. *Public Reason* **3**, 13–43.

Heath J. 2023. Anodyne privatization. *Erasmus Journal for Philosophy and Economics* **16**, 25–65. https://doi.org/10.23941/ejpe.v16i2.800.

Heilmann K., M.E. Kahn and C.K. Tang 2019. The urban crime and heat gradient in high and low poverty areas. *Journal of Urban Economics* 197. https://doi.org/10.1016/j.jpubeco.2021.104408.

https://doi.org/10.1017/S0266267124000282 Published online by Cambridge University Press

Howes A. 2022. Why innovation prizes fail. *Works in Progress* 7. https://worksinprogress.co/issue/why-innovation-prizes-fail/.

Kahn M.E., K. Mohaddes, R.N.C. Ng, M.H. Pesaran, M. Raissi and J. Yang 2021. Long-term macroeconomic effects of climate change: a cross-country analysis. *Energy Economics* 104. https://doi.org/10.1016/j.eneco.2021.105624.

Kranz G. 2000. *Failure is Not an Option: Mission Control from Mercury to Apollo 13 and Beyond*. New York: Simon and Schuster.

Kremer M. 2000a. Creating markets for new vaccines. Part I: Rationale. *Innovation Policy and the Economy* **1**, 35–72.

Kremer M. 2000b. Creating markets for new vaccines. Part II: Design issues. *Innovation Policy and the Economy* **1**, 73–118. doi: 10.1086/ipe.1.25056142.

Kremer M. and R. Glennerster 2004. *Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases*. Princeton: Princeton University Press.

Kremer M., J. Levin and C.M. Snyder 2020. Advance market commitments: insights from theory and experience. *AEA Papers and Proceedings* **110**, 269–273.

Kremer M., J. Levin and C.M. Snyder 2022. Designing advance market commitments for new vaccines. *Management Science* **68**, 4786–4814.

Kremer M. and H. Williams 2010. Incentivizing innovation: adding to the tool kit. *Innovation Policy and the Economy* **10**, 1–17.

Krugman P. 2011. An insurance company with an army. *New York Times*. https://archive.nytimes.com/krugman.blogs.nytimes.com/2011/04/27/an-insurance-company-with-an-army/.

Levine R., M. Kremer and A. Albright 2005. Making markets for vaccines. *Report of the Center for Global Development. Washington DC*. www.med.upenn.edu/istar/PDF-literature-photos/Making Markets-complete.pdf.

Malhotra A. and T. Schmidt 2020. Accelerating low-carbon innovation. *Joule* **4**, 2259–2267.

NASA 2014. *Commercial Orbital Transportation Services: A New Era in Spaceflight*.

NASA Office of Inspector General 2020. NASA's management of space launch system costs and contracts. https://oig.nasa.gov/docs/IG-20-012.pdf.

NASA Office of Inspector General 2021a. NASA's management of the Artemis Missions. https://oig.nasa.gov/docs/IG-22-003.pdf.

NASA Office of Inspector General 2021b. Artemis status update. https://oig.nasa.gov/docs/IG-21-018.pdf.

NASA Office of Inspector General 2022. 2022 Report on NASA's top management and performance challenges. https://oigforms.nasa.gov/docs/MC-2022.pdf.

NASA Office of Inspector General 2023a. NASA's transition of the space launch system to a commercial services contract. https://oig.nasa.gov/docs/IG-24-001.pdf.

NASA Office of Inspector General 2023b. NASA's management of the Artemis supply chain. https://oig.nasa.gov/docs/IG-24-003.pdf.

NASA Office of Inspector General 2023c. NASA's management of the space launch system booster and engine contracts. https://oig.nasa.gov/docs/IG-23-015.pdf.

Nemet G. 2019. *How Solar Energy Became Cheap: A Model for Low-Carbon Innovation*. London: Routledge.

Nocera J. and B. McLean 2023. Operation Warp Speed: the untold story of the COVID-19 vaccine. *Vanity Fair*, November 2023. www.vanityfair.com/news/2023/10/operation-warp-speed-covid-19-vaccine.

Park R.J., J. Goodman, M. Hurwitz and J. Smith 2020. Heat and learning. *American Economic Journal: Economic Policy* **12**, 306–339.

Pattison J. 2010. Deeper objections to the privatisation of military force. *Journal of Political Philosophy* **18**, 425–447.

Perrow C. 1999. *Normal Accidents*. Princeton: Princeton University Press.

Piacentini J., H. Frazis, P.B. Meyer, M. Schultz and L. Sveikauskas 2022. The impact of COVID-19 on labor markets and inequality. Bureau of Labor Statistics Working Paper 551. www.bls.gov/osmr/research-papers/2022/ec220060.htm.

Pinto S. 2023. The pandemic's effect on children's education. *Federal Reserve Bank of Richmond Economic Brief* 23–29. www.richmondfed.org/publications/research/economic_brief/2023/eb_23-29.

Roberts T.G. 2022. Cost for space launch to low-earth orbit: how much does it cost? Center for Strategic and International Studies Aerospace Security Project. https://aerospace.csis.org/data/space-launch-to-low-earth-orbit-how-much-does-it-cost.

Shklar J.N. 1989. The liberalism of fear. In *Liberalism and the Moral Life*, ed. N. Rosenblum, 21–38. Cambridge, MA: Harvard University Press.

Shklar J.N. 2023. Rights in the liberal tradition. *Political Studies* 71. https://doi.org/10.1177/00323217231171298.

Slaoui M. and M. Hepburn 2020. Developing safe and effective COVID vaccines – Operation Warp Speed's strategy and approach. *New England Journal of Medicine* 383. https://doi.org/10.1056/NEJMp2027405.

Smith A. 1981 [1776]. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Indianapolis: Liberty Fund.

United States Department of Health and Human Services 2020. U.S. Government engages Pfizer to produce millions of doses of COVID-19 vaccine. www.hhs.gov/about/news/2020/07/22/us-government-engages-pfizer-produce-millions-doses-covid-19-vaccine.html.

Vanderschraaf P. 2023. *Bargaining Theory*. Cambridge: Cambridge University Press.

Young I.M. 2011. *Responsibility for Justice*. Princeton: Princeton University Press.

Kirun Sankaran is Postdoctoral Fellow in the Political Economy Project and Lecturer in Philosophy at Dartmouth College. His research interests include the relationship between liberal political philosophy and innovation policy, social network methods, methods in political philosophy, and the intersection of political philosophy and the philosophy of science. URL: https://sites.google.com/view/kirunsankaran

Cite this article: Sankaran K. The body politic has private parts: market creation as a policymaking tool. *Economics and Philosophy*. https://doi.org/10.1017/S0266267124000282