



Mechanism design for pandemics

Eric Maskin¹

Received: 6 October 2021 / Accepted: 9 October 2021

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2021

Abstract

Under normal circumstances, competitive markets do an excellent job of supplying the goods that members of society want and need. But in an emergency like a pandemic, unassisted markets may not suffice. Imagine, for example, that society suddenly needs to obtain tens (or even hundreds) of millions of COVID-19 virus test kits a week. Test kits for this virus are a new product, and so it may not even be clear who the relevant set of manufacturers are. If we had the luxury of time, a laissez-faire market might identify these manufacturers: the price of test kits would adjust until supply matched demand. But getting a new market of this size to equilibrate quickly is unrealistic. Furthermore, markets don't work well when there are concentrations of power on either the buying or selling side, as there might well be here. Finally, a test is, in part, a *public* good (its benefits go not just to the person being tested, but everyone he might come in contact with), and markets do not usually provide public goods adequately. Fortunately, mechanism design can be enlisted to help.

Keywords Mechanism design · Pandemics · Markets

JEL Classification D71 · D79

1 Markets

Before getting to mechanism design, let's review why markets usually work so well. Suppose that there are many buyers and producers for some good. Suppose that buyer i enjoys (gross) benefit $b_i(x_i)$ from quantity x_i . Similarly, each producer j incurs cost $c_j(y_j)$ to produce y_j . Hence, society's net social benefit is:

$$\sum_i b_i(x_i) - \sum_j c_j(y_j) \quad (1)$$

✉ Eric Maskin
emaskin@fas.harvard.edu

¹ Harvard University, Cambridge, USA

At a social optimum, (1) is maximized subject to the constraint that supply equals demand:

$$\sum x_i = \sum y_j. \quad (2)$$

The solution to this constrained maximization is optimal in several senses:

- (i) total production $\sum y_j$ and total consumption $\sum x_i$ are optimal
- (ii) y_j is optimal for each producer j
- (iii) x_i is optimal for each buyer i

Achieving all three optimalities may seem complicated, but the market provides a simple solution. If p is the price at which the good can be bought and sold, then each buyer i maximizes

$$b_i(x_i) - px_i \text{ (net benefit)} \quad (3)$$

and the first-order condition for this maximization is

$$b'_i(x_i) = p \text{ (} b' \text{ denotes the derivative of } b \text{)} \quad (4)$$

Similarly, each producer j maximizes

$$py_j - c_j(y_j) \text{ (profit)} \quad (5)$$

with first-order condition

$$p = c'_j(y_j) \quad (6)$$

But notice that (4) and (6) are also the first-order conditions for the problem of maximizing (1) subject to (2). And so the market outcome attains the social optimum as long as p is chosen so that (2) holds (mathematically, p is the Lagrange multiplier for (2)).

How do we get the right choice of p ? In a free market, p falls if supply exceeds demand and rises if demand exceeds supply. Through an equilibration process, the right price is found, and so, as claimed, markets do indeed attain the social optimum.

But what if there is no pre-existing market for the good in question, as in the case of COVID test kits? We could still leave production and distribution to the free market, but there are several problems with doing so. In particular, how is a supplier to know (at least at first) how many tests kids to produce? After all, this is a new good and demand for it is uncertain. Furthermore, the supplier doesn't yet know who else will be producing test kits and how much they will produce. Under such circumstances, the supplier may be reluctant to incur the significant setup costs entailed in production until the uncertainties are resolved. Given time, the market could be expected to resolve them through the equilibration of supply and demand. But that process isn't instantaneous, and test kits are needed quickly.

Furthermore, given that supply can't be ramped up immediately, prices are likely to be high at first, which will disproportionately hurt poorer citizens and businesses (the very groups worst hit by the pandemic).

And finally, the market approach ignores the public good aspect of test kits. If I buy and use a test kit, I will get some benefit—I will know whether or not I have the virus and can take proper precautions and seek treatment if I do. But much of the benefit goes to *other* people, who will be protected from infection if I quarantine as a result of testing positive. Since I have little incentive to take into account those other benefits. I am likely to *underpurchase* test kits. And the market system will result in too few kits being supplied and used.

2 Mechanism design to the rescue

For all these reasons, we now turn mechanism design.¹ For now, let us assume that the government attaches (gross) benefit $b(\sum y_j)$ to total production $\sum y_j$ (in Sect. 4 we decompose $b(\sum y_j)$ into its component parts).

The government is interested in maximizing net social benefit

$$b\left(\sum_j y_j\right) - \sum_j c_j(y_j),$$

but it doesn't know the cost functions $\{c_j\}$ (and may not even know the full set of potential producers). We solve this difficulty using a variant of the Vickrey–Clarke–Groves mechanism (Vickrey 1961; Clarke 1971; Groves 1973). Specifically, the government announces a call for test-kit production and invites all potential producers j to submit a cost function \hat{c}_j . It then computes the production levels $\{\hat{y}_j\}$ that maximize the *apparent* net social benefit

$$b\left(\sum_j \hat{y}_j\right) - \sum_j \hat{c}_j(\hat{y}_j), \tag{7}$$

has each producer k produce \hat{y}_k , and gives producer k a payment

$$\left[b\left(\sum_j \hat{y}_j\right) - \sum_{j \neq k} \hat{c}_j(\hat{y}_j) \right] - \left[b\left(\sum_{j \neq k} \hat{y}_j^*\right) - \sum_{j \neq k} \hat{c}_j(\hat{y}_j^*) \right], \tag{8}$$

¹ An alternative to markets or mechanism design would be for government to simply *order* some company or companies to produce test kits. But this might be an extraordinarily inefficient outcome if these companies aren't up to the task or if there are other companies who could produce them much more cheaply (which the government is not likely to know in advance). Moreover, how does the government know what production level is "right"?

where the levels $\{\hat{y}_j^*\}_{j \neq k}$ maximize $b(\sum_{j \neq k} y_j) - \sum_{j \neq k} \hat{c}_j(y_j)$. The government then turns around and resells the test kits to citizens and others for a low (or even zero) price.

Claim Given that the government chooses $\{\hat{y}_j\}$ to maximize (7) and pays producer k the amount (8), it is optimal for producer k to report its costs *truthfully*, i.e., it will take $\hat{c}_k = c_k$.

Proof The second expression in square brackets in (8) doesn't depend on \hat{c}_k and so doesn't affect producer k 's maximization. In effect, producer k maximizes

$$b\left(\sum_j \hat{y}_j\right) - \sum_{j \neq k} \hat{c}_j(\hat{y}_j) - c_k(\hat{y}_k). \quad (9)$$

But (9) is just net social benefit with cost functions c_k and $\{\hat{c}_j\}_{j \neq k}$, i.e., producer k 's objective is the same as society's. Thus, the optimal choice of report \hat{c}_k is indeed c_k . \square

3 Simple example

Imagine that there is just a single producer with cost function $c(\cdot)$. In that case, the government

- (i) has the producer report $\hat{c}(\cdot)$
- (ii) calculates y^* to maximize $b(y) - \hat{c}(y)$
- (iii) has the producer produce y^*
- (iv) pays the producer $b(y^*)$
- (v) resells y^* to society

The producer's objective function is

$$b(y) - \hat{c}(y)$$

and so it is optimal for the producer to report $\hat{c} = c$.

Of course, there is a serious problem in this example: the supplier gets the *entire* social surplus. One way of avoiding this extreme outcome is to modify the government's payment to the supplier and make it

$$b(y) - k,$$

where k is a positive constant. This clearly doesn't affect the supplier's incentive to report its cost function truthfully. But it may be hard for the government to determine how big to choose k . If k is too small, the supplier still gets most of the surplus. If k is too big, the supplier may lose money, and so—without coercion—may refuse to participate.

However, the problem can typically be overcome by competition. If producer k is paid according to (8), then it is paid only its *marginal* contribution, i.e., how much it increases net surplus above what would occur in the supplier’s absence. Suppose for example, that other suppliers could supply the entire market efficiently (as would be the case if costs were linear in output). Then supplier k will only recoup its costs; it will earn no net rent (profit) at all. In other words, competition in this mechanism—in the free market—can be very effective at reducing suppliers’ rents.

4 Buyers’ benefits

Now, in general, $b(\cdot)$ might consist of public benefit $b_p(\cdot)$ and, for each consumer i , private benefit $b_i(\cdot)$. That is, $b(\{x_i\}) = b_p(\sum_i x_i) + \sum_i b_i(x_i)$, where the government has a good idea of b_p but not of the b_i s.

In this case, the government can have consumers report $\{\hat{b}_i\}$ (as well as having producers report $\{\hat{c}_j\}$) and, instead of maximizing (7), it can choose $\{\hat{x}_i\}$ and $\{\hat{y}_j\}$ to maximize

$$b_p\left(\sum_i \hat{x}_i\right) + \sum_i \hat{b}_i(\hat{x}_i) - \sum_j \hat{c}_j(\hat{y}_j) \text{ subject to } \sum_i \hat{x}_i = \sum_j \hat{y}_j \quad (10)$$

Consumer h then receives \hat{x}_h and pays

$$\left[\sum_j \hat{c}_j(\hat{y}_j) - b_p\left(\sum_i \hat{x}_i\right) - \sum_{i \neq h} \hat{b}_i(\hat{x}_i) \right] - \left[\sum_j \hat{c}_j(\hat{y}_j^*) - b_p\left(\sum_{i \neq h} \hat{x}_i^*\right) - \sum_{i \neq h} \hat{b}_i(\hat{x}_i^*) \right],$$

where $\{\hat{x}_i^*\}$ and $\{\hat{y}_j^*\}$ maximize

$$b_p\left(\sum_{i \neq h} x_i\right) + \sum_{i \neq h} \hat{b}_i(x_i) - \sum_j \hat{c}_j(y_j) \text{ subject to } \sum_{i \neq h} x_i = \sum_j y_j$$

By analogy with producer k ’s problem in Sect. 3, it is optimal for buyer h in these circumstances to set $\hat{b}_h = b_h$.

References

Clarke E (1971) Multipart pricing of public goods. *Public Choice* 11:17–33
 Groves T (1973) Incentives in teams. *Econometrica* 41(4):617–631
 Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J Finance* 16(1):8–37

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.