

Chapter 8 - Recognising Public Value: The Challenge of Measuring Performance In Government

Professor Mark Moore, Director, Hauser Center for Nonprofit Organizations, John F. Kennedy School of Government, Harvard University

Lecture presented 9 November 2006

Introduction

When I wrote *Creating Public Value*, my intention was to offer a challenge and a source of inspiration for people who wanted to provide leadership in the public sector. At the time, there was, quite rightly, some criticism along the lines of: 'Well, fine. One can readily talk about a concept of public value at a very high level of abstraction. But can we define the concept more concretely so that we can recognise objectively and empirically whether we are producing it?' This seemed like an important challenge to address.

The Concept of Recognising Public Value

To respond to this criticism, I began to talk about the challenge of '*recognising* public value' as well as *imagining* it and working to *produce* it. The word *recognising* seemed particularly suitable in talking about the challenge of measuring public value because it carried two different connotations – one intuitive, and one more technical.

The intuitive idea behind recognising public value is that we can collectively recognise a particular event or result as publicly valuable because we share a common intuitive notion of what we mean by public value. This is the meaning of 'recognition' suggested by United States Supreme Court Justice, Potter Stewart, who, when asked to define pornography, said, 'I know it when I see it'. So, one idea of recognising public value suggests some emergent, communal social process through which a community comes to understand (perhaps through some kind of continuing dialogue with an organisation) what valuable effects an organisation produces. Where such a common view exists, an organisation can take advantage of the shared understanding of the public value it delivers and simply point to those effects, as in: 'Look at that! We reduced crime in your neighbourhood,' or, 'See? The garbage has been removed from the curb'.

The second idea of recognising public value is far more technical. It derives from the field of financial accounting. The most basic requirement for financial accounting is the development of an administrative system that allows the organisation's managers to recognise expenditures of the organisation's assets and the receipt of financial assets by the organisation). These technical, administrative systems are constructed to render objective and consistent the movement of material and financial assets across the boundaries of an organisation and enable the use of those assets in productive activity. While standard financial accounting systems can capture financial flows through an organisation together with the costs expended by an organisation in producing particular products and services, it should be clear that if we want to measure the public value an organisation produces, we will have to construct some other technical system that can allow us to record when public value is being produced. Generally speaking, that requires us to develop some conceptual and operational basis for defining what constitutes the public value the organisation produces.

I will begin, then, by introducing the concept of 'recognising public value,' in both its abstract sense and its concrete manifestations.

The Example of COMPSTAT

One frequently cited example of the importance of developing the systems through which a community and an organisation can come to recognise public value comes from the New York City Police Department (NYPD). In January, 1994 Mayor Rudolph Giuliani announced the appointment of William Bratton as head of the NYPD.¹ Bratton came to the job saying he wanted to bring bottom-line management techniques to the NYPD. Correspondingly, like the CEO of a firm seeking to attract investment, he declared, first to the mayor and then to the public, that he was going to reduce crime by 10 per cent in New York City over the course of the next year. Somewhat surprisingly, this was a kind of promise that no police commissioner had ever made before.

Having promised this reduction, he exposed himself and his organisation to a substantial risk of failure. Necessarily, then, he began looking around for ways to engage other people in sharing the burden. He found a way to distribute responsibility for reducing crime throughout the NYPD by delegating responsibility for reducing crime to each of the NYPD's 72 precinct commanders. He and his deputy commissioners then called those managers to account for their crime reducing performance in public meetings in which precinct commanders were castigated for their failures to perform and challenged to come up with new approaches, or praised for strong performance with respect to the goal of reducing crime. They measured crime control effectiveness by looking at trends in reported crime statistics disaggregated by precinct and crime type. Essentially, Bratton followed the basic principles of 'Management 101': define your operational goals, delegate responsibility for achieving them, set up a system

for monitoring the performance of managers in achieving the desired results, and sanction the managers on the basis of their performance.

Amazingly, after establishing this system and laying out some new guidelines for patrol officers, reported crime rates fell substantially in the city. In fact, during the two-year tenure of Bratton and his management team, the murder rate in New York City dropped 50 per cent and the overall crime rate dropped 39 per cent. So, there was a rather astonishing moment when it looked as though we might be able to get significantly improved performance from police departments, or maybe even public sector organisations more generally, just by being sufficiently tough-minded and straightforward about measuring the performance of organisations with respect to their core mission. And thanks in part to this, the use of performance measurement as part of a performance management system in public sector organisations is now fundamental to the mainstream teaching of public administration.

Three Questions about Performance Measurement in the Public Sector

Despite the apparent success of this initiative, and despite the widespread commitment to these simple uses of performance measurement, I want to move the discussion of performance measurement up a level or two, because I do not think the simple lessons of this case are the right ones to draw. I am quite confident that the way that Bratton chose to measure the public value created by a police department is wrong, and wrong in ways that will apply to many more public agencies than police departments. And I am also reasonably confident that the method he chose for using performance measurement to animate and guide his organisation is not necessarily the best. But it is hard to argue with success – particularly when I am basically in agreement that we need more and better ways to measure the performance of public agencies. What I propose to do, then, is to back up a few steps from these simple lessons and ask three harder questions about performance measurement in the public sector.

The first is: Why should a public manager measure performance? Perhaps another way to put the question is, 'Why should public managers load a gun that others will use to shoot them?' I regard this basic motivational question as very important for two simple reasons. On one hand, building a strong performance measurement system requires a huge investment. On the other, having such a system inevitably exposes a manager and his organisation to the risk of a clearly documented failure. Given these facts, any moderately self-protective manager might reasonably ask, 'Why should I make this big effort when all it does is carry risk for me and for my organisation?' Even if all the rest of us citizens thought public managers ought to measure performance as part of their general accountability to us, it would help if we could find a compelling reason for managers to commit themselves to *doing* rather than *resisting* it.

The second question is: Where, along what I call the 'value chain,' is it best to measure performance? More particularly, the question might be, 'Should we measure inputs, processes, outputs or outcomes?'

The third question is: Should we measure customer satisfaction as an important dimension of public value creation, and, if so, who are the customers and what do they want?

Two Academic Traditions in Performance Measurement

I make these remarks against the backdrop of two substantial academic literatures that offer guidance about the recognition of public value on one hand, and the use of performance measures to animate and guide organisational performance on the other. The first is the academic literature that focuses on 'program evaluation' and 'benefit/cost analysis.' The second is the professional literature developed in business schools about the best ways to measure organisational performance.

Benefit/Cost Analysis and Program Evaluation

When I began my academic career over 30 years ago at the Kennedy School of Government, I was attracted, in part, because I thought that by going to the Kennedy School, I would learn a set of techniques that would allow me to undertake a relatively objective evaluation of the performance of government organisations. I thought that the development of those techniques was essential to improving both the design and evaluation of public policy initiatives. While I understood that there would inevitably be problems in achieving this goal, I certainly thought that by this time, almost thirty years later, we would have 'cracked' the issue of how to measure performance in the public sector.

At the Kennedy School, I was taught two basic ideas to use in recognising the public value of public policy initiatives and the performance of government organisations. The first of these was the theory and practice of 'benefit/cost analysis.' Benefit/cost analysis had a strong theoretical basis in welfare economics. The core idea was that individuals in a society were the only ones who could decide whether public policy results were valuable or not, and that the only way to discover what things were valuable was to ask individuals how they felt about them. The only way to find out the overall value of a given public policy was to sum up those individual valuations. If, for any given policy, there was enough value created for the beneficiaries of the policy to compensate those who suffered losses from it, then, arguably, some public value would be created by the proposed policy.

Although benefit/cost analysis had a strong theoretical basis, it was often applied in a less rarefied and more practical form. In that form, a benefit was defined as 'any effect that we consider good,' and a cost as 'any effect that we consider

bad.’ Thus, benefit/cost analysis was sometimes bowdlerised and used simply to identify the good things that were expected to come from a public policy (the benefits – whether individually experienced and evaluated or not), and the bad things that could happen (the costs – whether individually experienced or not.

The cruder form of benefit/cost analysis was closely related to a second method I was taught to recognise public value production. The second idea was called ‘program evaluation.’ Whereas benefit/cost analysis was grounded mainly in the discipline of economics, program evaluation came mostly from the field of statistics. The techniques of program evaluation required an analyst to develop a conception of public value creation by imagining what policymakers were trying to achieve through particular policies and programs. It was the hopes and aspirations of policy-makers to achieve particular social results that was the proper basis for defining value – not necessarily those of individuals who experienced the results (though we might have hoped that the policy-makers were guided by their best guesses about the views of individuals affected by public policies.) Once an analyst understood the initial aims and objectives of the policy or program, it would be possible to capture objective statistics to discern whether the characteristics of the world policy-makers were trying to influence did, in fact, change in the desired directions. One could then make a series of efforts to attribute observed changes to the activities carried out by the organisation. If the program evaluation demonstrated that the desired results had been achieved (at a low cost, with few unexpected bad consequences) then we could conclude that public value had been created. If the program evaluation showed that the desired results had not been achieved, or that the costs or unexpected adverse consequences had been larger than anticipated, then we could conclude that public value had been lost. This was how, as a student in public policy, I was taught to measure the performance of organisations. I subsequently discovered that this was a difficult, somewhat clumsy, and very expensive approach.

Performance Measurement in the Private Sector

I also learned that in the private sector, businesses did not think about performance measurement in this way at all. Rather, managers in the private sector used their organisation’s information systems to track whether the organisation was creating value in the form of revenues over costs, or whether they were managing to maintain a satisfactory relationship with their customers. They were also, surprisingly, making much use of something that looked very similar to the traditional methods we were using in the public sector, which was essentially the measurement of concrete processes and activities. These kinds of measures helped organisations in both sectors make sure they were getting the benefit of technological systems.

Case studies of the McDonald's and Burger King restaurant chains published by the Harvard Business School revealed a surprising fact. I naturally assumed that McDonald's and Burger King would analyse the performance of their restaurants in financial terms. Much to my astonishment, I learned that they paid little attention to the financial performance of each individual restaurant. Instead, McDonald's was evaluating the performance of every restaurant on 12 distinct characteristics of the *experience* of being in the restaurant, including whether it was clean; whether the lines were short; whether the rest rooms were clean; whether there was a place to sit down; whether the food was fresh; and whether the cashier smiled at you when she took your money.

These concrete aspects of the experience of being in the restaurant, rather than hamburgers and fries, was McDonald's theory of what they were producing. What they wanted from each of their restaurants was a guarantee that each of the 12 attributes of the experience was being produced with a high degree of consistency. Their theory was that if they produced consistently good performance along the 12 dimensions of customer experience, they would make money. That theory was tested quarterly, and annually, against the financial performance of the organisation as a whole. But that theory didn't have to be tested for each individual restaurant. For the most part, all the individual restaurants needed to do was manage themselves to produce the particular set of 12 product/service attributes.

What was really eye-opening, however, was how often they measured the performance of each restaurant: about once every month. Over the period of a year, each restaurant would have been evaluated 12 times on each of these criteria. If we were to suggest some similar kind of performance measurement regime to a public sector executive, the response would likely be, 'Oh my God, that is an incredibly expensive and intrusive information system!' Bear in mind that McDonald's is collecting this information even when they have at hand a steady flow of financial information on the performance of the organisation.

One possible lesson for public sector executives is that if a firm like McDonald's is making this many extra measurements even when they have financial data available, then it might make sense for them to spend at least that much on measurement when they don't have good financial information. A second possible lesson is that measurement carried out this frequently might allow top managers to detect relatively small changes in organisational performance – changes that managers could conceivably influence. The fact that small changes in performance can actually be detected gives this performance measurement system more behavioural power to shape the performance of midlevel managers and staff than a less rigorous system would have.

An interesting question arises about which of these two different ways of thinking about these issues – the program evaluation/benefit-cost perspective

or the business management perspective – offers a more viable approach to measuring performance in the public sector.

Question 1: Why Measure Performance? Strategic Uses of Performance Measurement

At this point, I will introduce some additional complications. In a narrow sense, there is a standard set of ideas about the managerial uses of performance measures. They are:

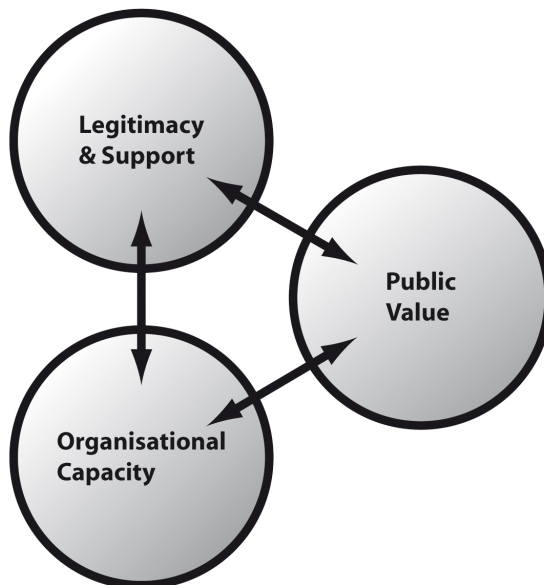
1. to meet demands for external accountability;
2. to establish a clear, significant mission and goal for the organisation; and
3. to foster a strong sense of internal accountability.

This last use of performance measurement could mean simply making everybody feel accountable for working harder to meet target numbers by setting and enforcing rigorous performance standards, but, if we capture information about both what we are doing and the results we get, could also allow for a softer kind of accountability admits the possibility of failure and allows more room for learning and innovation.

The Strategic Triangle

As many of you know, in *Creating Public Value* I developed a framework for thinking about strategic management in government that could be reduced to and represented by a ‘strategic triangle.’ (Figure 1)

Figure 1 - The Strategic Triangle



I argued that it was the important goal of public sector executives to find a 'fit' between a conception of the 'public value' their organisation sought to produce, and ensuring a continuing flow of resources to the organisation by building 'legitimacy and support' in what I call the 'authorising environment.' In addition, for any particular conception of public value to be strategically viable, managers would have to be able to build sufficient operational and organisational capabilities to achieve results. What was strategic about this was the idea that we were trying to fit a public sector organisation – led by an executive with control over a bundle of assets and a bundle of capabilities – into a dynamic environment. This environment is, on the one hand, a 'task environment' in which we try to orchestrate a portfolio of programs, policies, and procedures to produce public value and, on the other hand, an 'authorising environment,' consisting of the citizens, legislators, interest groups, and other parties that tell the manager what particular pieces of the world that executive is authorised and expected to attend to.

When I thought about the standard arguments for performance measurement in the public sector against the backdrop of the strategic triangle, I realised that performance measurement – the process of defining and recognising public value production – fits rather neatly into the strategic triangle and can be used as a highly effective, even essential, tool for strategic management. Take each point in order.

First, at the legitimacy and support point of the triangle, it is, I think, easy to see that one of the best ways to guarantee our legitimacy – to invite support – is to embrace accountability for defining and recognising public value. Indeed, we have a professional and ethical obligation to be accountable for our performance to those who have the power to call us to account and expect us to perform, and this means developing concrete methods of measuring specific aspects of performance.

Second, at the public value point of the triangle, the effort to develop and use performance measures forces us to be much more concrete and explicit about the nature of the public value we are trying to create on behalf of our authorisers. If we cannot make it concrete, we cannot measure it.

Third, at the operational capacity point of the triangle, it is essential that we develop the administrative wherewithal to get the most capacity out of our organisations. This makes it necessary to foster a strong sense of internal accountability, and to give those who work in the organisation guidance so that they can focus on finding the means of improving performance through organisational learning and reflection. These responsibilities depend on the capacity to measure the performance of the organisation concretely and reliably.

So, for a manager standing in the middle of the strategic triangle, it is not hard to see that performance measurement would play an absolutely critical role with

respect to handling each corner of this triangle: the definition of public purpose, the mobilisation of support, and the exploitation of operational capacities by driving performance inside the organisation (whether through a rigid structure of accountability or through a softer accountability that emphasises learning).

Negotiating the Terms of Accountability

Performance measurement also plays a key role in bringing the different points of the strategic triangle into alignment. Indeed, managers can use performance measurement to shape their mandate for action. They do this by trying to negotiate the terms of accountability with their overseers in the authorising environment. Given the importance of having a steady and coherent mandate for action in managing their organisations, it might well be worth a public manager's time and effort to use the challenge of developing performance measures for his or her organisation as a key part of developing their overall strategy.

If we, in our roles as citizen authorisers in democratic countries, were consistent in our expectations of public sector managers, and managers could expect the authorising environment to make sustained and serious efforts to demand performance, then measuring performance would be a very important feature of their task environment. Public managers might feel the need to invest significant effort in negotiations with the political authorising environment to reach some kind of agreement about what they are expected to produce. If public managers were trying to make an agreement with politicians about what they ought to produce and measure, it would be important to them that the politicians not be permitted to change the rules suddenly. If one could get an agreement to stick for a period of time, one would have, in effect, created a policy mandate and, if the authorising environment remained stable in its expectations and demands, managers would at least have some reasonable chance of accomplishing their agreed-upon goals.

Unfortunately, while the political system often says that it wants accountability, and occasionally demands it, it is usually reluctant to negotiate a stable deal about what constitutes public value. There is too much political advantage in retaining the right to change the rules when they want and finding the organisation suddenly accountable for some value that had not previously been treated as important. The reality is that the political system will not allow itself to be negotiated into a binding, long-term agreement. It will reserve the right to change the rules whenever it wants. If we cannot negotiate with the political authorising environment to achieve an agreement about performance measurement, then negotiating in advance will not help us manage organisational performance.

Doing ‘The Right Thing’

Perhaps, the right thing to do, then, is to appeal to shared virtues of executive responsibility and commitment. To be perfectly candid, I think this is an important part of the duty of a public official and, frankly, I cannot imagine any conscientious public official feeling fully accountable and fully responsible in his job without taking some responsibility for defining whatever it is that he or she is trying to accomplish and making some effort to measure successes and failures. I do not think people can manage public organisations without having a conception of public value that can be measured reliably. Performance measurement is fundamental to the ability to manage organisations. If you do not have such a tool, then I would argue that you are merely *presiding* over the organisation’s activities. You are not really *managing* it and you are not creating conditions under which the organisation’s performance could be expected to improve in any reliable way over a period of time. And the reason I think this is so important is that, left to their own devices, organisations tend to accumulate slack, and the only way to move that slack around and get it performing is to embrace a strong form of external accountability in order to animate and drive performance within the organisation.


This is exactly what Bill Bratton did in New York City. When Bratton said, ‘I am going to reduce crime by 10 per cent over the next year,’ he embraced external accountability, set a high standard, and exposed himself and his organisation to the risk of failure. This approach produces a certain amount of energy inside an organisation.

Now, as it turns out, this approach could also be unhelpful in the sense that it might not be a constructive challenge, or in the sense that it might drive the organisation in the wrong direction. However, it is also the case that, if we start running a little bit of voltage through the organisation – not too little, not too much – something begins to happen inside the organisation. The reason I think one should measure performance is that I do not believe one can manage or lead an organisation without some kind of accountability – that is to say, without concrete performance measurement.

Question 2: Where Should We Measure Performance Along the Value Chain?

A second key question in designing performance measurement systems focuses on the best place to measure performance in the complex system that transforms inputs granted to the organisation into the production of desired results through a particular set of policies, programs, and activities. One useful way to think about this production process is to describe it as a value chain that consists of several distinct stages of the production of public value. (See Figure 2 below)

Figure 2 - The 'Value Chain'



Obviously, the concept of a value chain is not rocket science. The value chain describes an organisation that is receiving *inputs* (in the case of a public organisation, money and authority from the state), which are deployed in particular *processes, procedures, and programs* which produce *outputs*. Many of those outputs involve an important engagement with people, or *clients* – and by clients I mean people who make individual transactions or have individual encounters or relationships with the organisation. One of the important consequences, or products, of those client transactions or encounters or relationships, is some form of *social outcome*. So the value chain gives us a general picture of organisational production.

But the value chain, as it is drawn above, also reminds us to focus on the actions of *partners and co-producers* who can be instrumental in the production of client satisfaction and social outcomes. Contributions from partners and co-producers could emerge spontaneously without any encouragement or guidance from the organisation. Or, they could emerge as a consequence of deliberate efforts made by the organisation to mobilise co-production activities. The organisation can write contracts with partners or it can seek to use moral persuasion of various kinds to motivate private individuals and organisations to contribute to public goals. It can even use the authority of the state to require others to contribute to socially desired outcomes.

In principle, we could take measurements and demand accountability for performance at any step along the value chain or the chain of partners and co-producers. We could focus on inputs to make sure that our managers have firm control over the equipment and inventory, that they are doing a good job of cash management, and so on. We could focus our attention on activities, processes, procedures and programs through a 'compliance audit,' where

independent evaluators would go into government organisations and see whether people were following the policy and procedures that were laid down for them.

Measuring Activities, Processes, and Programs

We can all see the limits of a compliance audit: it does not tell us very much about the quantity or quality of outputs we are producing, or whether the outputs are judged to be valuable in the sense of satisfying clients or producing desired social outcomes. However, if we had already embedded in our policies and procedures the best available technology for accomplishing results, then a compliance audit would be tantamount to an audit of efficiency and effectiveness. If our procedures were known to be efficient and effective in producing desired results, then all we would have to do is monitor and demand performance with respect to following policy and procedures and, arguably, we would have achieved our purpose. If the policies and procedures are based upon 'best practices,' we can effectively guarantee operational performance by adhering to these policies and procedures rather than by guaranteeing results. (This is similar to the approach that McDonald's takes in trying to ensure that each restaurant produces the desired products and services. The judgment about whether this particular product and service can be delivered in a profitable way is made elsewhere.)

But there is another reason to adhere to strict policy and procedural regimes: namely, to assure not efficiency and effectiveness, but *consistency*. This is a desirable characteristic of government operations quite apart from questions of efficiency and effectiveness. One of the characteristics we want in a government organisation is a guarantee that like cases will be treated alike. This is part of what we mean by fairness, and it constitutes a virtue independent of how effective a particular government process turns out to be.

So, our options are: (1) concentrate on managing performance measurement at the level of activity, processes and procedures; or (2) monitor outputs, workload measures, and measures of productivity. Increasingly, however, government organisations are being pushed to measure performance even further down the value chain in the direction of either client satisfaction or outcomes. I think the reasons for that are pretty clear. But, as you know, the great debate goes on about whether we should measure outcomes or outputs in the public sector.

Measuring Social Outcomes

My observation is that the current debate is shifting heavily in the direction of measuring outcomes. The reason for this is that an outcome is a direct measure of the public value that we are trying to produce. If we need to be able to demonstrate that we are producing something of public value, then we would have to be able to measure outcomes. Once we devise a way to measure the outcome we are trying to achieve, we can find out whether the technical means

that we are relying on represent the best technology for accomplishing that purpose, or whether there is a better alternative we can use to test the value proposition offered by the organisation, so that when we say, 'Well, we are going to produce public value by engaging in the following activities,' we can track whether our theory is correct.

If it is so clear that outcome measurement is valuable, why do we so rarely do it? Why do we not measure outcomes? Well, it turns out there are a lot of reasons why measuring outcomes is a difficult device for managing organisations. One reason is that measuring outcomes is incredibly expensive.

In a private sector firm, revenue is earned by the sale of products and services. Revenues earned by the sale of products and services to willing customers are easy pieces of information to collect. They are a terrific thing to use in measuring performance because they represent a reasonably objective, direct measure of value. What makes them a direct measure of value? Individuals' choices to spend their own money on a given product or service: people plunk their money down to buy the product or service. In doing so, they give objective information that they value whatever good or service the firm is producing.

The second attractive property of revenue earned through the sale of products and services (as an outcome measurement) is that the information is collected right at the boundary of the organisation at the moment of transaction. We don't have to go beyond this transaction to measure something that takes place at a relatively remote place and time. Nor do we have to wait for the result to occur. We get the judgment of value right at the point of delivery.

A third property of revenues earned through sales is that one can compare apples and oranges. Individuals show us how much they value these distinct products by paying for them in a given currency.

So, earning a revenue from the sale of products and services turns out to be a really wonderful way to measure the value of organisational output. If we have such a measurement, managers (and others) are in a good position to judge the value of what has been produced. If we do not have revenues earned from the sale of products and services to voluntary customers, we are in big trouble.

To see the significance of the loss of this kind of revenue measure, consider this problem: if I were the manager of General Motors and somebody said to me: 'You can have all the information you want about the cost of producing automobiles, but you cannot have any information on how much you have earned by selling them,' what would I do? How would I know whether I was creating something valuable?

I think there are only a limited number of answers. One of them would be to survey your customers to see whether they liked the vehicle. Another possibility: you might observe the outcome and see how often people drive and how well

the car stands up to use. A third possibility: you might try to get engineers to tell you whether they think the car that you designed is a good car – is it better than previous or comparable models?

These would be all the same methods we would use in the public sector to try to ascertain whether we are producing something of value. We survey our customers; we evaluate outcomes; we check on the physical characteristics of our product. While such measures might do some managerial work for us in helping us determine whether we are producing valuable results, none of these measures of value would stack up well against the information we would receive from measuring the revenue that we earned from the sales of products and services.

So, one of the difficulties with relying on the measurement of outcomes is that outcomes are very expensive to measure: we have to go out into the world to find the information and, often, information comes in late.

In the case of a welfare-to-work program, for example, we are interested in trying to spend public dollars to cause clients who use the program to move from a dependent life on welfare to an independent status of job-holding and an active life in the community. An interesting question is: When do we decide that we have accomplished that outcome? And the answer to this question is: We never *have* achieved the outcome – at least not for a large group of people. What we are trying to do is to maintain a certain level of functioning for a group of people who, ideally, rely less and less on public support in the form of both dollar benefits and social services. The benefit of public support comes from the fact that people improve in their independence and social functioning continuously over a long period of time. So, when do we decide that we have accomplished the goal?

Perhaps the best way to think about it is that, each month, we have kept X number of people at a higher level of functioning than they would have been without intervention. That number represents the value but, in order to get it, you have to have continuous monitoring of the whole target population over a long period of time – a very difficult and expensive process.

Further complicating things, because the information about outcomes comes in late and public value production happens far down a causal chain, the effect is often hard to attribute to public managers or public organisations. Crime statistics may have gone down under Bill Bratton's watch, but we actually do not know whether or to what degree that was the consequence of any given action that Bratton took.

While we talk a great deal about the ultimate importance of focusing on outcomes in the public sector, and while the challenge to think about and measure the ultimate outcomes is certainly valuable, the reality is that we are not only

interested in the social outcomes produced by government organisations. We are also interested in the processes and procedures we use to achieve the outcomes. We want to make sure they have certain characteristics – including fairness and economic and technical efficiency – and that means that we will want to measure some characteristics of activities and outputs as well as outcomes. So, while I understand and support the argument for measuring outcomes, there are real problems with relying exclusively on outcome measures.

Measuring Outputs

So then, what about measuring outputs? Measuring outputs is cheaper and easier than measuring outcomes. In addition, measuring outputs can contribute to improving the performance of midlevel managers because, unlike outcomes, outputs are mostly within managers' direct control. We can make interesting calculations about productivity and – without necessarily making claims about value or effectiveness – we can at least see whether we are driving down the cost per unit of output. This helps to focus the manager's attention. We get improved consistency in operations, and, if we have an idea about what kinds of practices constitute value, we can make sure that those best practices are being used widely throughout our organisation.

These observations add up to some important reasons to measure outputs rather than outcomes. So much so that it might be tempting to rely only on the measurement of outputs. The problem with that solution, however, is that organisational outputs are not quite the same as the production of something of public value – not even in terms of producing client satisfaction, let alone in terms of the achievement of desired social outcomes.

If both outcomes and outputs are important and useful to measure, and each has limitations, it seems to me that the right conclusion is that one ought to measure both.

This conclusion is reinforced by the observation that private sector organisations are increasingly managing according to 12 or 13 or 14 different measures of performance. They do not limit themselves to one 'bottom line.' They exist in a 'data rich environment' where they can make judgments about the degree to which their current strategies are succeeding and, at the same time, search for clues about how they might improve their performance.

The bad news about measuring multiple dimensions of performance is that it makes it harder for managers to focus. The more things you are trying to manage, the more complicated the job becomes. Another way to think about it, however, is that if we are managing 12 different dimensions of performance, then we have an opportunity to engage the organisation in a conversation about what it is being produced. A conversation about how to get all of those different dimensions of performance moving in the right direction would be more interesting for

those inside the organisation than a simple thumbs up or thumbs down test of how well the organisation and the individuals in it are performing. The organisation might actually be able to learn and feel motivated to find better ways of achieving difficult results, rather than simply censuring people for failing to meet more or less arbitrary standards.

Question 3: Should We Measure Customer Satisfaction? Who are the Customers and What Do They Want?

The last question about performance measurement is the degree to which we should measure customer or client satisfaction, and who, exactly, is the customer? I have been struggling for a long time with this idea of a customer-oriented government. Personally, I do not favour the concept of a customer-oriented government because I do not actually think the government has ‘customers.’ Rather, I think the government has *clients* on one hand, and *citizens* on the other.

The characteristic of a customer in the private sector is that the person engages in an individual financial transaction with the private sector organisation. Customers are *practically* important to private sector firms because they provide financial support to the organisation through the voluntary purchase of goods and services. Insofar as a private sector firm can produce products and services that appeal to customers, the private firm can stay in business and continue to accumulate wealth for its shareholders.

But customers are *normatively* important to private sector firms as well. It is the customers’ individual, voluntary decisions to purchase goods and services that provide the basic moral and social justification for the enterprise. They are, in essence, the arbiters of the value created by the enterprise. If customers can be satisfied (without deception or coercion), then one can make the argument that society as a whole ought to be satisfied with the output of the private firms.

Now, I am a public sector guy and proud of it. But every now and then, my strong public sector orientation causes me to say something really stupid. I get on my high horse, and I say something like, ‘Why does the private sector produce junk like lemon-scented furniture polish and hula-hoops? With all the problems we have in the world, why on earth are we spending any managerial time, any labour, any resources on things like lemon-scented furniture polish and hula-hoops?’ I feel really good when I say that. The problem is that the private sector has a really good answer: ‘People *buy* this junk. The fact that they are buying it tells us that they value it. The fact that they value these products at a price that more than covers the cost of production means not only that businesses can stay alive and profitable, but also that they are creating something of value – as long as we think it is individuals’ judgments of value, revealed in decisions to spend their own money, that defines at least one kind of value.’

This is why customers are really important, both practically and ethically, in the private sector.

Customers as Clients: Beneficiaries vs. 'Obligatees'

So, who are the customers in the public sector? If we focus on the question of the relationship between the person and the organisation, we might imagine that the customer in the public sector would be the *client* of the organisation: the person who engages in a transaction at the business end of the organisation, rather than the reporting end. We might think this is a particularly apt definition in the case of clients who are beneficiaries of the organisation. Thus, it is easy to imagine that welfare recipients, or schoolchildren, or people who call the police for service – all of whom are transacting with the organisation at the individual level and getting something beneficial – might be thought of as customers of public sector organisations.

Notice, however, that there are some *clients* of government organisations who do not receive beneficial services from government. These clients include criminals who are arrested by the police, taxpayers who are asked to fork over their hard-earned money to the government, and companies who are required to stop dumping toxic wastes in the air and water. Instead of receiving beneficial services from government organisations, these particular clients of the organisation receive an obligation. I call those clients 'obligatees'.

This term amounts to a kind of linguistic crime. I know this is true because whenever I type 'obligatees' on my computer, a squiggly little red line shows up: this clearly is not a word in the English language. However, I am going to persist in using the word because I think we need it to understand how to manage government organisations. It conveys the essence of an important relationship between the government and individuals in society.

Government is in the business of obliging individual clients to contribute to public purposes as well as providing services to clients. Consequently, we have to think about what these kinds of clients would value in what I would call an 'obligation encounter' as well as a service encounter. The interesting question becomes: What are the socially and individually desired properties of obligation encounters?

I would go so far as to say that most, maybe even all, individual transactions between government organisations and individual clients involve both a service component and an obligation component. When we give money to welfare recipients, or teach students in public schools, we are not simply providing a service whose value will be judged by the clients themselves. In both cases, we are delivering a service that comes with certain social expectations and aspirations, and the quality of that service will be judged not only by the

recipients, but also by the citizens and taxpayers whose aspirations to achieve some socially valued condition are contained in these transactions.

If we (tentatively) accept the claim that many clients of public services are indeed 'obligatees', then it becomes apparent, I think, that the goal of these transactions could not be simply to satisfy these clients and make them happy. Given a free choice, I expect most obligatees would say, 'If you really wanted to make me happy, you could excuse me from the obligation you are about to impose'. If we accept the claim that the sole objective of an obligation encounter cannot be to make an obligatee happy, then government must be producing these obligation encounters for some other purpose. That purpose might lie further down the value chain as a social outcome associated with the results of an obligatee coming into compliance with the duty imposed upon him or her.

Customers as Citizen Authorisers

An alternative idea about the 'customers' of government services is that they are not the clients who meet organisations in individual transactions at the service end of the organisation, but rather that they are the citizens of the democracy who authorise and pay for the organisations and hope that the organisations will realise through their operations the aggregate social results that citizens want to achieve. They are bystanders and witnesses of governmental organisations, rather than the objects of the organisations' attention. Their aspirations about what the government organisation ought to do in general become the guide and arbiter of the value produced by the organisation. It is citizens – acting collectively and deciding that some particular purpose (a social outcome) is worth taxing and regulating themselves to accomplish – who provide the practical and normative basis for a government organisation to proceed.

The first thing I notice, if we define the customers of government as citizens rather than clients, however, is that they do not look very much like customers. If we insist on using private sector analogies, they look more like shareholders – whose interests are reflected in the judgments of a board of directors – than customers.

In a sense, these citizen overseers and authorisers are more fundamentally the organisation's 'customers' than the people who interact with the organisation through individual transactions. After all, it is the citizens who provide the financial support to the organisation. It is also the citizens' collective decisions to support this or that public enterprise that constitute the normative justification for the organisation to continue its operations. In the public sector, the arbiter of public value is not just the individuals who transact with the organisation, it is the collective acting as a whole to achieve collectively defined social goals. The difficulty, of course, is that citizens are a collective customer – not simply a bundle of individual customers. Consequently, they have to learn how to speak

with a clear voice through the mechanisms of representative government about what exactly they would like to have produced to meet their collective needs or aspirations.

So, let us compare and contrast *customers* in the private sector with *clients* and *citizens* in the public sector. In the private sector, customers meet businesses in individual transactions and encounters. In this sense, they resemble clients of public sector organisations. In the private sector, customers pay a price that covers the full cost of producing the goods and services. In the public sector clients do not – even when they are receiving a service from government. In addition, many clients in the public sector get obligations as well as services.

In the private sector, customer satisfaction is an important business goal. In the public sector, client satisfaction may be an important goal, but it is not necessarily the most important goal. In the private sector, customers' individual decisions to purchase a good or service for their own use generally provides both the practical and normative basis for the private sector firm to continue. In the public sector, clients' decisions about whether or not they like the service they receive does not play this important role. Instead, the resources available to a public organisation come from a collective decision made by citizens and their representatives that the purposes to be pursued by a given public organisation are valuable enough to tax and regulate themselves to accomplish. In the public sector, it is citizens – not customers – who provide the resources an organisation needs to carry on, and who act as the final arbiter of the value produced by the organisation.

Reasons to Measure Client Satisfaction

Of course, there are still some reasons why individual client satisfaction in encounters with government agencies could play an important role in a public enterprise. Public sector clients have rights that need to be protected in the delivery of services. To the extent that the individuals confronted by government organisations enjoy having their rights protected, we would want to make a commitment to satisfying individual clients along this particular dimension of performance. We, as citizens, could also decide and declare that an important role of the public organisation is to make clients happy and to treat them with dignity. If we, acting collectively as citizens rather than as individual clients, made this declaration, then the organisation would have to treat the satisfaction of individual clients as an important goal (among others, such as the achievement of social outcomes at a low cost). Or, if making clients happy helped to achieve something that citizens have declared a socially valuable result, then that, too, would be a reason to want to keep clients happy.

But the point is that when it comes to evaluating the performance of government organisations, we cannot assume that the goal is to make clients happy. We have

to explore the relationship between client satisfaction and the social goals that citizens, acting through the representative processes of government, have set for the organisation.

I draw the conclusion, then, that customer satisfaction in both service encounters and obligation encounters in the public sector should be designed not only to protect rights and ensure some level of client satisfaction but also to produce any changes necessary to achieve desired outcomes. This makes it necessary to distinguish between those outcomes desired by the collective client 'WE' and those desired by the individual client 'I'. In other words, as citizens acting through collective political processes, we might articulate desired outcomes for individual clients, and, ideally, we are also collectively trying to help those individual clients help us achieve shared purposes.

Giving Citizens What They Want

As I have noted, the authorisation and tax revenues provided by citizens keep public sector organisations alive. But another important question for public managers to ask is: What do citizens, acting in their role as arbiters of value, want in exchange for their tax dollars and their liberty? Let us consider the possibilities. The simplest idea of what citizens might want from public organisations might be social outcomes. A slightly more complicated idea might be social outcomes at a relatively low cost in terms of the use of money and authority. You might add client satisfaction. There are various possibilities.

The point is that the organisation should be oriented to the citizen authorisers and their aspirations rather than to individual clients and their level of satisfaction. Presumably, citizens acting collectively might want various attributes of performance from public organisations. In policing, for example, citizens might want reduced crime, reduced fear, a certain kind of justice associated with calling offenders to account for their crimes, another kind of justice associated with the protection of individual rights, or still another kind of justice associated with a fair distribution of police services across a population – all at the lowest possible cost. Indeed, it is often true that citizens demand attributes of performance that seem inconsistent with one another, or unreasonable.

The important question then becomes: How do we cope with this incoherent mandate? One answer is simply to do the best we can to give those in the authorising environment whatever seem to want from us. The answer could be that we provide to each constituency, shareholder, or authoriser a report that speaks to their concerns. Furthermore, we ought to define value creation in the organisation in terms of movements in desired directions along as many attributes of performance as there are ideas about the value created. The idea here is that we should think of the performance of an organisation not in terms of a single measure of performance, but instead as a bundle of performance attributes, each

of which is an object of concern to some kind of citizen, or some group that is participating in the process of democratic oversight.

An initial objection to such an approach would be to claim that it would be impossible for an organisation to improve its performance in all dimensions simultaneously. Surely, we cannot have less crime and reduced costs of policing! Surely we can't have greater reliability in both calling offenders to account and protecting civil rights! We need guidance about how to make trade-offs among competing values, not simply instruction about which of many dimensions of performance have value.

Public Sector Organisations and the Production Possibility Frontier

It is worth remembering, however, that an organisation will be forced to make trade-offs among competing performance goals only if it is now operating on what economists would call the production possibility frontier – the point at which an organisation can be said to be making the best possible use of its assets to produce socially valuable results.² It is only when we have exhausted all possible technical improvements in our efforts to produce social output that we necessarily face trade-offs. If we are not at the production possibility frontier, then, in principle, it is possible for organisations to discover technical means that will allow them to improve performance on some attributes of performance without sacrificing performance on other (apparently inconsistent) attributes of performance. In short, there may be different methods or technologies of policing that would allow us to have less crime *and* lower costs; or to catch more offenders without violating citizens' rights.

If trade-offs among competing values are necessary only when organisations are at their production possibility frontier, an important question is: What is the likelihood that public sector organisations are operating at the production possibility frontier? My answer is that it is very *unlikely* that most public sector organisations are even close to their theoretically possible production possibility frontiers. The reason is simply that public organisations face fewer pressures and have fewer opportunities to search for their production possibility frontiers than private sector firms. And even with all the pressures and opportunities that push private sector firms face to the production possibility frontier, many private sector firms still never get there.

Indeed, the fact that many private sector organisations are not on the production possibility frontier was revealed several decades ago when American automobile manufacturers faced a stiff competitive challenge from Japanese automakers. The Japanese were exporting cars to the United States that were better than American cars in all dimensions of quality. Further, they were \$2,000 cheaper than anything produced in the US. Faced with this competitive challenge,

American car manufacturers thought they were stuck. They thought they could not increase quality without increasing costs, and that they were beaten on both dimensions. In short, they thought they were on their production possibility frontier. It turned out that they were wrong. Once they began looking at their production processes in search of technical improvements, they found many were possible. With investments in new operating technologies, they found they could simultaneously improve quality and reduce costs.

The lesson for the public sector might be that we should not assume at the outset that our public organisations are already operating at their production possibility frontier. In fact, we ought to make the opposite assumption; that it might be possible to improve along all dimensions of performance. We should not shift to the assumption that we face trade-offs until the evidence generated by performance measurement systems that capture all relevant dimensions of performance tell us we are now facing important trade-offs among values. In short, it might be wise to try to improve performance on all dimensions of performance simultaneously. That would help us discover how close we were to production possibility frontiers as we try to make good on our commitments to produce the things that our authorisers value.

Improving Along Multiple Dimensions of Value: An Example from Policing

One of the most important things in management is simply to know whether we are improving, not whether we have found the optimal solution. It is improvement that we are after. The optimal solution is going to change tomorrow when conditions and aspirations change. So, in the short run, the important thing to do is to focus on improvement, and improvement can best be guided by a comprehensive list of the attributes of performance that citizens and their representatives deem valuable in public sector operations.

Going back to my friend Bill Bratton, he formulated the proposition that the 'profit' the NYPD earned for New York City's citizens was reduced crime. Now Bratton was a very good manager, but a miserable accountant. When you think about it, the idea that reduced crime actually equals profit is not quite right. In any standard accounting system, profit equals revenues minus costs. When trying to estimate the profit earned by the police, then, we would have to take into account the costs (in both dollars and sacrifices of personal liberty) of producing reductions in crime. In addition, we might have to recognise that reduced crime is not the only thing that citizens value in a police department. There might be a large number of valuable effects other than crime control that the police manage to produce.

Over the years, I have spent a great deal of time thinking about how to measure public value creation in policing. I came to the conclusion that crime reduction

was better seen as the equivalent of the 'gross revenue' earned by the NYPD for the citizens of New York City, not the profit. The profit (or net public value) earned by the NYPD would be the value citizens attached to reduced crime minus the costs to citizens of producing that result.

Policing relies on inputs of both public funds and public authority. Precisely because the police use authority as well as money (and are engaged in obligation as well as service encounters), police operations have to be evaluated in two quite distinct normative frameworks: a *utilitarian* framework to assess the effectiveness of inputs of labour and resources on one hand, and a *justice* framework to assess the degree to which the police helped to produce justice and behaved justly and fairly in the pursuit of both the practical goals of reducing crime and the principled goal of assuring justice. In the utilitarian framework, the end is to reduce crime and the means are to spend resources on materials to produce that result. In the justice framework, the end is to produce justice, and the means include using the force and authority of the state. A comprehensive evaluation would have to focus on whether both money and authority were used efficiently, justly, and fairly in efforts to achieve the results.

After a great deal of reflection and discussion, I hypothesised that we could actually understand most of the public value we citizens want the police to produce in terms of just seven different dimensions of performance. These were as follows:

1. Reducing crime and victimisation;
2. Producing the kind of justice associated with practical redress and calling offenders to account;
3. Reducing fear and creating a sense of personal security;
4. Ensuring stability in public spaces and creating ordered liberty;
5. Using force and authority fairly, efficiently, and effectively;
6. Using financial resources fairly, efficiently, and effectively; and
7. Trying to make people who call for services happy with the service they receive.

To test the adequacy of this conception, it seemed important to follow the idea that a good performance measurement system was one that spoke to the real concerns of citizens and other authorisers. Consequently, my colleagues and I carried out empirical investigations in two cities to determine what citizens and authorisers wanted from their police departments. This required us to identify both those who were actively overseeing police operations and the particular dimensions of performance that were the focus of their concern.

The important authorisers included four broad categories:

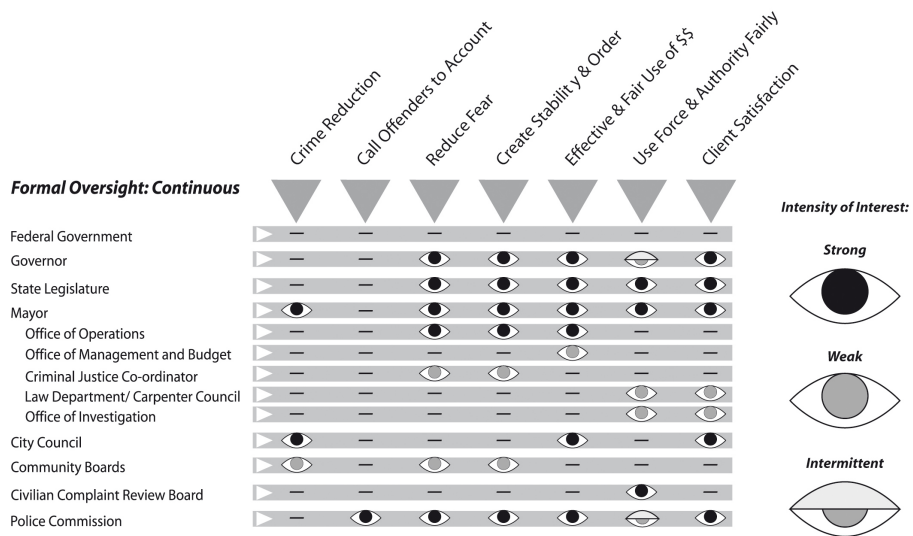
1. Formal overseers exercising continuous monitoring of organisational performance;

2. Specially commissioned oversight agencies;
3. Informal organisations that appointed themselves to oversee the police; and
4. Partners in the criminal justice system.

These are the people who oversee, judge, mandate and monitor the value of police performance.

We then looked to see what attributes of performance interested them. Figure 3 presents the results obtained for New York City. Figure 3 sets out the particular dimensions of police performance particular ‘authorisers’ in NYC were paying attention to, and indicates the level and consistency of attention each authoriser is giving to a particular dimension of performance.

Figure 3 - What Authorisers of Police Care About



The results presented in this figure suggest, surprisingly, that relatively few people turned out to be worried about ‘reducing crime.’ An awful lot of attention seems to be focused on increasing a sense of security at the lowest possible cost. This gives both citizens and police managers a different basis on which to consider the social goals of policing, particularly when we understand both that fear is not always highly correlated with real criminal victimisation, and that some citizens are at least as afraid of the police as they are of muggers.

It is always possible that managers of public sector organisations might hold views about which of their organisation’s activities create public value that are fundamentally different from those of authorisers and/or clients. This raises a strategic question for managers about how best to develop an external constituency that shares their conception of the public value they are trying to get their organisation to produce. One way for managers to search for such a

constituency would be to construct a measurement system to monitor a particular activity or feature of performance they think is valuable to see who, if anybody, pays attention to it. This measurement system could become a tool for managing the organisation as well as a tool for delivering the activity in question.

Conclusion

In conclusion, performance measurement plays an essential role in creating public value through effective strategic management. The work of developing and improving performance measurement systems involves philosophical and normative as well as scientific and cognitive issues. Every time I make a claim that something is valuable, and therefore worth measuring, I make a philosophical/normative claim, not just an empirical/positive claim. Performance measurement is about value. Necessarily, then, performance measurement has important political dimensions beyond its obvious administrative and technical dimensions. This has to be the case because the only place a manager can go to get guidance about public value is the political authorising environment.

The obvious way for a manager to initiate and carry on a dialogue with the political world is to ask those in the political authorising environment what they think is valuable. If managers can engage their political overseers in serious, sustained discussions about the definition of public value, or the publicly valued dimensions of performance for their organisations, then they might be able to come to some conclusions about the definition of public value. Engaging in such a discussion, however, might well expose the manager and his organisation to criticism, and to the risk of a real, objective failure to produce what the public wants.

There are only two things that might motivate a manager to begin working in this risky and murky terrain – to take up the challenge of defining and recognising the creation of public value with the input and oversight of the political authorising environment. The first is that by doing so, one might emerge with a stronger, clearer, more consistent definition of the public value one is being asked to produce. The second is that one might call into existence the conditions under which one can successfully manage and lead one's organisation.

What finally compels managers to construct the systems that will allow them to recognise public value is that it is impossible to live up to the duties of their offices if they do not know what they are trying to produce in sufficiently concrete terms to know whether they are succeeding or not, and if they do not have the measurement tools that allow them to drive performance and seek out the technical means for continuing improvement. Only the strategic use of performance measurement makes such things possible. However, engaging in the task of constructing performance measures always means confronting unresolved conflicts. In the end, it means exposing ourselves and our

organisations to potential failure. In the face of that, we might fall back in alarm. The only thing that moves us forward is the knowledge that we cannot run organisations without really understanding what constitutes public value and how we contribute to its creation.

Bibliography and Notes

Mark H. Moore (1995), *Creating Public Value: Strategic Management in Government*, Harvard University Press.

ENDNOTES

¹ In 1994, William Bratton was appointed the 38th Commissioner of the New York City Police Department by Mayor Rudolph Giuliani. He introduced the CompStat system of tracking crimes, which proved successful in reducing crime in New York City and is still used to this day, but left the job in 1996 after alleged personal conflicts with Giuliani. In 1996, Bratton was featured in a business case prepared by James L. Heskett and published by Harvard Business School. Bratton's efforts to effectively turn around the New York City Police Department is used by many business schools, including Kenan-Flagler Business School, as a tool for teaching organisational design and change. (Source: http://en.wikipedia.org/wiki/William_J._Bratton)

² The production possibility frontier (PPF) represents the point at which an economy is most efficiently producing its goods and services and, therefore, allocating its resources in the best way possible. If the economy is not producing the quantities indicated by the PPF, resources are being managed inefficiently and the production of society will dwindle. The production possibility frontier shows there are limits to production, so an economy, to achieve efficiency, must decide what combination of goods and services can be produced. (Source: <http://www.investopedia.com/university/economics/economics2.asp>)