

Policy Retrospectives

Douglas J. Besharov
Editor

Submissions to Policy Retrospectives should be sent to Douglas J. Besharov, School of Public Policy, University of Maryland/American Enterprise Institute, College Park, Maryland 20742; (301) 405-6341.

Incentives and Their Dynamics in Public Sector Performance Management Systems

Carolyn J. Heinrich and Gerald Marschke

Abstract

We use the principal-agent model as a focal theoretical frame for synthesizing what we know, both theoretically and empirically, about the design and dynamics of the implementation of performance management systems in the public sector. In this context, we review the growing body of evidence about how performance measurement and incentive systems function in practice and how individuals and organizations respond and adapt to them over time, drawing primarily on examples from performance measurement systems in public education and social welfare programs. We also describe a dynamic framework for performance measurement systems that takes into account strategic behavior of individuals over time, learning about production functions and individual responses, accountability pressures, and the use of information about the relationship of measured performance to value added. Implications are discussed and recommendations derived for improving public sector performance measurement systems. © 2010 by the Association for Public Policy Analysis and Management.

INTRODUCTION

The use of formal performance measures based on explicit and objectively defined criteria and metrics has long been a fundamental component of both public- and private-sector incentive systems. The typical early performance measurement system was based largely on scientific management principles (Taylor, 1911) promoting the careful analysis of workers' effort, tasks, work arrangements, and output,

establishing work procedures according to a technical logic and setting standards and production controls to maximize efficiency. Assuming the benchmark level of performance reflected the relationship between worker effort and output, workers could be paid according to a simple formula that included a base wage per hour plus a bonus rate for performance above the standard.

The basic concepts underlying this simple compensation model—that employees perform better when their compensation is more tightly linked to their effort or outputs, and organizational performance will improve with employee incentives more closely aligned with organizational goals—have since been broadly applied in the design of private- and public-sector performance measurement systems. Richard Rothstein (2008) presents a wide-ranging, cross-national, and historical review of such systems, including Soviet performance targets and incentives for enterprise production; performance indicators in the U.S. Centers for Medicaid and Medicare Services and Britain's National Health Service; incentive systems for bus drivers in Santiago, Chile; and the growing use of performance accountability systems in public schools.

Bertelli and Lynn (2006, pp. 28–29) suggest that the early technical focus of public sector performance measurement systems was congruent with a prevailing “scientism in political science,” orienting toward descriptive analysis of formal governance structures and processes rather than attention to the dynamics of system incentives and their consequences. Almost a century later, public managers still have relatively little leeway for basing employees' pay on their performance, although the U.S. Civil Service Reform Act of 1978 sought to remedy this by allowing performance-contingent pay and increasing managers' latitude for rewarding performance. The ensuing development and testing of “pay for performance” systems for both personnel management and public accountability is ongoing, albeit research suggests that public-sector applications have to date met with limited success, primarily due to inadequate performance evaluation methods and underfunding of data management systems and rewards for performance (Rainey, 2006; Heinrich, 2007).

Over time, researchers and practitioners have come to more fully appreciate that the conditions and assumptions under which the simple, rational model for a performance measurement and incentive system model works—organizational goals and production tasks are known, employee efforts and performance are verifiable, performance information is effectively communicated, and there are a relatively small number of variables for managers to control—are stringent, and in the public sector, rarely observed in practice. The fact that many private- and public-sector production technologies typically involve complex and nonmanual work, multiple principals and group interactions, political and environmental influences and interdependencies, and nonstandardized outputs makes the accurate measurement of performance and construction of performance benchmarks (through experimental studies, observational methods, or other approaches) both challenging and costly. As Dixit (2002) observed, these richer “real world” circumstances often make it inappropriate to rely heavily on incentives linked in simple ways to output measures that inexactly approximate government performance. Yet, because of their ease of use and the significant costs associated with developing and implementing more sophisticated measures and more intricate contracts or systems of incentives, simple incentives built around inexact and incomplete measures are still widely used in performance management systems for inducing work and responsible management.

Recent public sector reforms aimed at improving government performance and the achievement of program outcomes likewise reflect this simplistic approach to motivating and measuring performance. The U.S. Government Performance and Results Act (GPRA) of 1993, along with the Program Assessment Rating Tool (PART) used by the Office of Management and Budget, require every federal agency, regardless of the goals or intricacy of their work, to establish performance goals and

measures and to provide evidence of their performance relative to targets in annual reports to the public. Furthermore, distinct from earlier public sector performance measurement systems are the focus on measuring *outcomes* and the greater demands for external visibility and accountability of these performance measurement activities, along with higher stakes attached to performance achievements through increasing use of performance-contingent funding or compensation, organization-wide performance bonuses or sanctions, and competitive performance-based contracting (Radin, 2006; Heinrich, 2007). In his comprehensive study of performance management reforms, Moynihan (2008, pp. 4–5) asks, “[h]ow important are these performance management reforms to the actual management of government?” He answers, “[i]t is only a slight exaggeration to say that we are betting the future of governance on the use of performance information.”

As a result of these ongoing efforts to design and implement performance management systems in complex organizational settings, we have been able to compile a growing body of information and empirical evidence about how these incentive systems function in practice and how individuals and organizations respond and adapt to them over time. The primary goals of this “Policy Retrospectives” piece are to synthesize what we know both theoretically and empirically about the design and dynamics of the implementation of public sector performance measurement systems and to distill important lessons that should inform performance measurement and incentive system design going forward. The research literature on performance measurement is diverse in both disciplinary (theoretical) and methodological approach. We therefore choose a focal theoretical frame—the principal–agent model—that we argue is the most widely applied across disciplines in studies of performance measurement systems, and we draw in other theories to elaborate on key issues. In addition, our review of the literature focuses primarily on studies that produce empirical evidence on the effects of performance measurement systems, and in particular, on systems that objectively measure performance outcomes—the most commonplace today (Radin, 2006). We discuss examples of performance management systems, largely in public social welfare programs (education, employment and training, welfare-to-work, and others), to highlight important challenges in the design and implementation of these systems, both in the context of a broader theoretical framework and in their application. We also acknowledge that there is a very rich vein of qualitative studies that incisively investigate many of these issues (see Radin, 2006; Frederickson & Frederickson, 2006; Moynihan, 2008), although it is beyond a manageable scope of this review to fully assess their contributions to our knowledge of performance management system design and implementation.

In the section that follows, we briefly describe alternative conceptual or theoretical approaches to the study of performance management systems and incentives and then focus the discussion on key elements of a commonly used static, principal–agent framework for informing the development of performance incentive systems. We also highlight the limitations of particular features of these models for guiding policy design and for our understanding of observed individual and organizational responses to performance measurement requirements. We next consider the implications of a dynamic framework for performance management system design that takes into account the strategic behavior of agents over time, learning about production functions and agent responses, and the use of new and/or better information about the relationship of measured performance to value added. We are also guided throughout this review by a basic policy question: Do performance management and incentive systems work as intended in the public sector to more effectively motivate agencies and employees and improve government outcomes? In the concluding section, we discuss recommendations for improving public sector performance management systems and for future research in this area.

THEORETICAL MODELS FOR PERFORMANCE MANAGEMENT AND INCENTIVE SYSTEMS

Theoretical models that have been applied in the study of performance management and incentives range from a very basic “logical model” of “relevant factors” in a performance management system (Hatry, 1999) to more elaborate economic and socio-psychological frameworks that consider multiple principals and levels of influence, complex work technologies and group interactions, intrinsic and extrinsic motivations, pecuniary and non-pecuniary rewards, political and environmental influences and interdependencies, and other factors. Hatry’s simple model links inputs, activities, outputs, and outcomes (intermediate and end) and describes relationships among these different types of performance information. Although his model does not formally identify the influence of context or relationships among performance measures at different levels, he does call for public managers to gather explanatory information along with their performance data—from qualitative assessments by program personnel to in-depth program evaluations that produce statistically reliable information—to interpret the data and identify problems and possible management or organizational responses.

In an approach that clearly diverges from the scientific logic, Moynihan (2008, p. 95) articulates an “interactive dialogue model of performance information use” that describes the assembly and use of performance information as a form of “social interaction.” The basic assumptions of his model—that performance information is ambiguous, subjective, and incomplete, and context, institutional affiliation, and individual beliefs will affect its use—challenge the rational suppositions underlying most performance management reforms to date. While performance management reforms have emphasized what he describes as the potential “instrumental benefits” of performance management, such as increased accountability and transparency of government outcomes, Moynihan suggests that elected officials and public managers have been more likely to realize the “symbolic benefits” of creating an impression that “government is being run in a rational, efficient and results-oriented manner” (p. 68). Moynihan argues convincingly for the cultivation and use of strategic planning, learning forums, dialogue routines, and related approaches to encourage an “interactive dialogue” that supports more effective use of performance information and organizational learning. He also asserts that performance management system designers and users need to change how they perceive successful use of performance information. Performance information systems and documentation of performance are not an “end” but rather a means for engaging in policy and management change.

Psychologists’ study of performance management and incentives has focused primarily on the role of external rewards or sanctions in shaping and reinforcing individual behaviors, particularly in understanding the relationship between rewards and intrinsic motivation (Skinner, 1938; Deci, Koestner, & Ryan, 1999). Organizational psychologists conducting research in this area have highlighted the role of employee perceptions about the extent to which increased effort will lead to increased performance and of the appropriateness, intensity, and value of rewards (monetary or other). Motivation theorists, alternatively, have stressed the importance of intrinsic and autonomous motivation in influencing individual behavior and performance, arguing that extrinsic rewards are likely to undermine intrinsic motivation. Meta-analyses by Deci, Koestner, and Ryan (2001) report mixed findings on the effects of performance incentives, depending on the scale of rewards and competition for them, the extent to which they limit autonomy, and the level of difficulty in meeting performance goals.

Studies of public-sector performance management and incentives have also drawn on stewardship theory, which traverses the fields of psychology and public management and assumes that some employees are not motivated by individual goals but align their interests and objectives with those of the organization (Davis, Donaldson, &

Schoorman, 1997; Van Slyke, 2007). In other words, organizational stewards perceive that their personal needs and interests are met by achieving the goals of the organization. Through this theoretical lens, individuals engaged in public service such as teaching are seen as less motivated by monetary incentives because they genuinely buy into the mission of public education and are intrinsically motivated to increase their students' knowledge. In the public administration literature, empirical tests of the theory of "public service motivation" provide support for the tenet that some public employees are motivated by "an ethic to serve the public," acting because of their commitment to the collective good rather than self-interest (Perry & Porter, 1982; Rainey, 1983; Perry & Wise, 1990; Houston, 2006). Implications for the design of incentive systems are that giving financial rewards to these employees may be counterproductive, as they may exert less effort depending on the alignment of their public service motives with performance goals (Burgess & Rato, 2003).

The principal-agent model that we focus on in this paper is not only one of the most widely applied in the study of performance management and incentive systems across disciplines, but it also yields important insights that are fundamental to the design and implementation of any incentive system. For example, it provides a cogent framework for addressing asymmetric information, one of the most common problems in assessing agent effort and performance and aligning the incentives of multiple system actors (organizations or individuals). The principal-agent model is also particularly informative for analyzing how organizational structures, performance measurement system features, and management strategies influence the responsiveness of agents to the resulting incentives. In the common applications of this model that we study, the agent typically operates the program (performs the tasks) that are defined by the principal to achieve organizational or program goals, although we acknowledge that with the increasing role of third-party agents in public services delivery, it is frequently the case that agents have only partial or indirect control.

A Static Framework for Performance Measurement and Incentive System Design

In outlining a static framework for the design of incentive systems in public organizations, we begin with the multitasking model (Holmstrom & Milgrom, 1991; Baker, 1992), typically presented with a risk neutral employer and a risk-averse worker.¹ An organization (the principal) hires a worker (the agent) to perform a set of tasks or actions. The organization cannot observe the worker's actions, nor can the organization precisely infer it from measures of organizational performance. We assume the goal of the principal is to maximize the total value of the organization. If the organization is a firm, total value is measured as the present value of the net income stream to the firm's owners. Alternatively, public-sector organizations maximize value added for the citizens they serve, net of the public's tax contributions to public services production. Organizational value can derive from multiple, possibly competing objectives. Value may derive not only from improvements in student academic achievement, in public safety or in individuals' earnings, but also from equitable access to services, accountability to elected officials, and responsiveness to

¹ The principal's risk preferences are assumed to be different from the agent's because the model is frequently used in contexts where the principal is the stockholder of a for-profit firm and the agent is a worker in the firm. The firm's stockholders may have a comparative advantage in risk management because they can more easily spread their wealth among different assets, while workers may have most of their wealth tied up in their (non-diversifiable) human capital, limiting their ability to manage their risk. Moreover, persons selecting into occupations by their ability to tolerate risk implies that the owners of firms will be less risk averse than the firms' workers. In the public sector, the goal of the organization is not profit but policy or program (possibly multidimensional) objectives, and the principal is not the owner of a firm but a public sector manager or a politician. The public sector principal may not as easily diversify away the risk of managerial or policy failures by purchasing assets whose risk will offset policy risk (a point made by Dixit, 2002, p. 5). Thus the assumption that the principal and agent have different risk attitudes in public sector applications of the model may not always be as justifiable.

citizens. The set of actions that a government employee can take to affect value will potentially be moderated by the effects of uncontrollable events or environmental factors that impact value (e.g., economic and political variables). The literature typically assumes the organization does not directly observe these actions or the random environmental factors that impinge on them, but it does know the functional relationship among the variables, including their marginal products and the distribution from which the random error (intervening, uncontrollable factors) is drawn.²

The assumptions that the principal cannot detect and therefore reward effort (or a given agent's contribution to productivity), and that effort is costly to the agent, suggest that the agent will exert a suboptimal level of effort or will exert the wrong levels of effort across productive tasks, if paid a flat salary.³ If value were observable, the organization could base the employee's compensation on his or her value added because it is directly affected by employee actions, thereby creating a clear incentive to exert effort. In the public sector, the value added of programs is occasionally assessed through the use of randomized experiments (such as MDRC's experimental evaluations of welfare reform strategies, including the New Hope Project, Minnesota Family Investment Program, and others), or as in education, through sophisticated statistical modeling of value added, but, more often than not, it is not observed.⁴ Alternatively, there are often one or more performance measures available that reveal information about the employee's actions, such as the test scores of students in public schools, which are assumed to reflect their teachers' instructional skills and school supports. Any performance measure, however, will also capture the effects of random, uncontrollable events or environmental factors on the performance outcome.⁵ As with value, it is assumed that the organization doesn't observe these factors but does know the functional relationship among the variables, the marginal product of employees' actions, and the distribution of the random error. Note that while they may share common components and may be correlated, the random factors associated with measuring value added and measuring performance are generally different.

Relationship Between Performance and Value Added and Choice of Performance Measures

The observation that the effect of an action on performance may be different than its effect on value—and that this difference should be considered in constructing a wage or work contract—is a key insight of the research of Baker (1992) and Holmstrom

² The relationship between value added and the government employee's actions is often modeled as a linear one. Following Baker's (2002) notation, let value added $V(\mathbf{a}, \varepsilon) = \mathbf{f}\mathbf{a} + \varepsilon$, where \mathbf{a} is a vector of actions that the worker can take to affect V , \mathbf{f} is the vector of associated marginal products, and ε is an error term (mean zero, variance σ_ε^2) that captures the effects of uncontrollable events or environmental factors.

³ What we mean by suboptimal allocation of effort is an allocation of effort that deviates from the allocation the principal would command if the principal were able to observe the agent's actions. The key assumption, and what makes for a principal-agent "problem," is that the agent's objectives are *different* from the principal's objectives. This assumption is consistent with a wide range of agent preferences, including those conforming to versions of the public service-motivated civil servant described in the public administration literature. It is incompatible with a pure stewardship model of government, however. If either the preferences of the agent are perfectly aligned with those of the principal or the agent's actions are verifiable, there is no role for incentive-backed performance measures in this framework.

⁴ To be clear, we are using the term "value" or "value added" in an impact evaluation sense, where the value added of a program or of an employee to a given organizational effort is assessed relative to the outcome in the absence of the program or the employees' efforts, or the counterfactual state that we do not observe.

⁵ Following Baker's (2002) notation, we model performance (P) as: $P(\mathbf{a}, \varphi) = \mathbf{g}\mathbf{a} + \varphi$, where φ is an error term (mean zero, variance σ_φ^2) that captures the effects of uncontrollable events or environmental factors. Note that \mathbf{g} is not necessarily equal to \mathbf{f} , because P and V are typically measured in different units, and, crucially, the sets of actions that influence P and V , and how they influence P and V , are generally different. Although the uncontrollable events or environmental factors that influence P are different from those that influence V , ε and φ may be correlated.

and Milgrom (1991). Both the theoretical and empirical literature also recognize the possibility that some actions that increase performance may simultaneously *reduce* value added. An example of an action that influences performance but not value is a teacher providing her students with (or filling in) the answers to a state-mandated test (Jacob & Levitt, 2003). Differences in the functions that generate performance and value may also induce the worker to exert the right kinds of effort but in the wrong quantities: a teacher “teaching to the test.”

By linking the worker’s compensation to a performance measure, the organization may increase the worker’s attention to actions that also increase value. One of the simplest compensation schemes is a linear contract: The agent receives a salary independent of the agent’s effort and performance, plus a sum that varies with performance. The variable part is a piece rate—a fixed amount of compensation per unit of performance. The higher the piece rate, the greater is the fraction of the agent’s compensation that stems from performance and the stronger the incentive to exert effort. On the other hand, the higher the piece rate, the greater the influence of factors outside the agent’s control on compensation, and thus, the greater the risk taken on by the agent.

Baker’s (1992) model shows how the organization should navigate this trade-off, which has been studied extensively in the literature (Holmstrom, 1982). The sensitivity of an agent’s compensation to performance should be higher the more beneficial (and feasible) an increase in effort is. Increasing the intensity of performance incentives is costly to the organization, because it increases the risk the agent bears, and correspondingly, the compensation necessary to retain her. Thus, raising the incentive intensity makes sense only if the added effort is consequential to the value of the organization. Everything else equal, the sensitivity of an agent’s compensation to performance should be lower as the agent’s risk aversion and noise in the measurement of outcomes increase. That is, the less informative measured performance is about the agent’s effort, the less the principal should rely on it as a signal of effort.

Incentives should also be more intense the more responsive the agent’s effort is to an increase in the intensity of incentives. That is, incentives should be used in organizations for agents who are able to respond to them. In some government bureaucracies, workers are highly constrained by procedural rules and regulations that allow little discretion. Alternatively, in organizational environments where agents have wide discretion over how they perform their work, agents may respond to performance incentives with innovative ways of generating value for the organization. Imposing performance incentives on agents who have little discretion or control needlessly subjects them to increased risk of lower compensation.

In addition, Baker’s model introduces another trade-off that organizations face, namely, the need to balance the effort elicited by a high piece rate (based on performance) with the possible distortions it may cause. The sensitivity of an agent’s compensation to performance should be smaller the more misaligned or distortive the performance measure is with respect to the value of the organization. A key implication of this model is that the optimal incentive weight does not depend on the sign and magnitude of the simple correlation between performance and value, which is influenced by the relationship between external factors that affect the performance measure and those that affect the organization’s value. Rather, what matters is the correlation between the effects of agent actions on outcomes and on the organization’s value.⁶

For example, from the perspective of the policymaker or taxpayer who would like to maximize the value from government dollars spent on public programs, we want

⁶ To sum up, in terms of the notation in footnotes 2 and 5, the optimal incentive weight falls with the risk aversion of the agent and the angle between \mathbf{f} and \mathbf{g} (the misalignment of P and V), and rises with the magnitude of \mathbf{g} relative to σ_ϵ^2 (the effort signal-to-noise ratio).

to choose performance measure(s) so that the effects of agent actions on measured performance are aligned with the effects of those same actions on value. For a given performance measure, we would like to understand how closely the performance and value functions align, but as we frequently don't have complete information about them, this is difficult to realize in practice. Empirical research on this issue has focused primarily on estimating measures of statistical association between performance and value, which may tell us little about the attributes of a given performance measure that make it useful, such as the degree of noise and its alignment with value.

The choice of performance measures is further complicated by the fact that organizations typically value performance across a set of tasks, and that the workers' efforts are substitutes across tasks. The optimal contract will then place weight on the performance measure corresponding to each task, and the agent will respond by exerting effort on each task (Holmstrom & Milgrom, 1991). As is often the case, however, suppose that a performance measure that is informative exists for only a subset of the tasks. Holmstrom and Milgrom show that if an organization wants a worker to devote effort to each of several tasks or goals, then each task or goal should earn the worker the same return to effort at the margin; otherwise, the worker will devote effort only to the task or goal that has the highest return to effort. The implication is that in the presence of productive activities for which effort cannot be measured (even imprecisely), weights on measurable performance should be set to zero. Or as Holmstrom and Milgrom (1991, p. 26) explained, "the desirability of providing incentives for any one activity decreases with the difficulty of measuring performance in any other activities that make competing demands on the agent's time and attention," which is perhaps why (at least in the past) the use of outcome-based performance measures in public social welfare programs was relatively rare.

Some researchers have suggested that in such organizational contexts, subjective performance evaluation by a worker's supervisor allows for a more nuanced and balanced appraisal of the worker's effort. In a recent study of teacher performance evaluation, Jacob and Lefgren (2005) compared principals' subjective assessments of teachers to evaluations based on teachers' education and experience (i.e., the traditional determinants of teacher compensation), as well as to value-added measures of teacher effectiveness. They estimated teachers' value added to student achievement using models that controlled for a range of student and classroom characteristics (including prior achievement) over multiple years to distinguish permanent teacher quality from the random or idiosyncratic class-year factors that influence student achievement gains. Jacob and Lefgren found that while the value-added measures performed best in predicting future student achievement, principals' subjective assessments of teachers predicted future student achievement significantly better than teacher experience, education, or their actual compensation. Two important caveats were that principals were less effective in distinguishing among teachers in the middle of the performance distribution, and they also systematically discriminated against male and untenured teachers in their assessments. Baker, Gibbons, and Murphy (1994) argue that a better result may be obtained if evaluation is based on both objective and subjective measures of performance, particularly if the two measures capture different aspects of employee effort.

Clearly, subjective evaluation is imperfect, and it may be more problematic than alternatives because it produces measures that outside parties cannot verify. Furthermore, subjective measures are also subject to distortion, including systematic discrimination such as that identified by Jacob and Lefgren (2005). At the same time, it has also been argued by Stone (1997, p. 177) and others (Moynihan, 2008) that *all* performance information is subjective; any number can have multiple interpretations depending on the political context, because they are "measures of human activities, made by human beings, and intended to influence human behavior."

Employee Motivations and Choice of Action

One might ask why employees would take actions that increase only measured performance, with little effect on value. Using public employment and training programs as an example, consider the technology of producing value, such as a real increase in individuals' skills levels, compared to what is required to arrange a job placement for a program participant (one of the long-used performance measures in these programs). Based on prior research in workforce development, (Barnow, 2000; Courty & Marschke, 2004; Heckman, Heinrich, & Smith, 2002), we suggest that it will require greater investments of resources and effort on the part of employees, such as more intensive client case management and skills training provided by more experienced program staff, to affect value compared to the level of effort required to increase performance. Given that value does not enter into the employee's compensation contract, the employee will have little incentive to choose actions that incur a higher cost to him (to produce value), particularly if these actions do not correlate strongly with actions that affect measured performance.

At the same time, as discussed earlier, there is a growing interdisciplinary literature that suggests that individuals in organizations may be motivated to exert effort in their work by something other than the monetary compensation they receive, such as by an ethic to serve the public or to act as a steward of their organization. The implications of intrinsic, stewardship, or public-service motivations among employees are that they derive utility (i.e., intrinsic rewards) from work and will exert more effort for the same (or a lower) level of monetary compensation; and through their identification with the goals of the principal or organization, they are also more likely to take actions that are favorable to the interests of the principal (Akerlof & Kranton, 2005; Murdock, 2002; Prendergast, 2007; Van Slyke, 2007; Francois & Vlassopoulos, 2008). In other words, if value added is the goal of the organization (e.g., increasing clients' skill levels in an employment and training agency) and is communicated as such by the principal, intrinsically motivated employees should be more likely to work harder (with less monitoring) toward this goal and should be less likely to take actions that only increase measured performance and more likely to take actions that only increase value. The returns on exerting effort in activities in the correct quantity to maximize value added are greater for the intrinsically motivated agent, and as he is less interested in monetary pay, he will be less concerned about increasing performance to maximize his bonus.

The findings of Dickinson et al. (1988), Heckman, Smith, and Taber (1996), Heckman, Heinrich, and Smith (1997), and Heinrich (1995) in their studies of public employment and training centers under the U.S. Job Training Partnership Act (JTPA) program aptly illustrate these principles. Heckman, Heinrich, and Smith (1997) showed that the hiring of caseworkers in Corpus Christi, Texas, who exhibited strong preferences for serving the disadvantaged, in line with a stated goal of JTPA to help those most in need, likely lowered service costs (i.e., wage costs) in Corpus Christi. Alternatively, an emphasis on client labor market outcomes in a Chicago area agency that was exceptional in its attention to and concern for meeting the federal performance standards appeared to temper these public-service motivations (Heinrich, 1995). With performance expectations strongly reinforced by administrators and in performance-based contracts, caseworkers' client intake and service assignment decisions were more likely to be made with attention to their effects on meeting performance targets and less so with concern for who would benefit most from the services (the other basic goal of the JTPA legislation).

In the undesirable case in which both principal and agent motivations and the effects of agent actions on performance and value do not align, the consequences for government performance can be disastrous, as Dias and Maynard-Moody (2007) showed. They studied workers in a for-profit subsidiary of a national marketing research firm that shifted into the business of providing welfare services. As they

explained, requirements for meeting contract performance (job placement) goals and profit quotas generated considerable tensions between managers and workers with different philosophies about the importance of meeting performance goals versus meeting client needs. They noted that the easiest way to meet contract goals and make a profit was to minimize the time and effort line staff devoted to each client, as individualized case management requires “detailed, emotional interaction with clients, taking time to discover their areas of interest and identifying the obstacles preventing them from accomplishing their individual goals,” which distracts from the goal of immediate job placement (p. 191). In the face of these pressures, managers had little incentive to go beyond minimal services, and they consequently assigned larger caseloads to frontline staff to allow less individual time with clients. As Dias and Maynard-Moody reported, however, the differences in views of the intrinsically motivated staff (focused on value) and managers concerned about performance triggered an “ideological war” that subsequently led to “enduring and debilitating organizational problems” and failure on both sides to achieve the program goals (p. 199).

In addition to monetary and intrinsic incentives, empirical research also suggests that agents respond strongly to non-pecuniary incentives, which, in the public sector, most frequently involve the potential for public recognition for performance (Bevan & Hood, 2006; Walker & Boyne, 2006; Heinrich, 2007). Incentive power is typically defined in monetary terms as the ratio of performance-contingent pay to fixed pay, with higher ratios of performance-contingent, monetary pay viewed as “higher-powered,” that is, offering stronger incentives (Asch, 2005). However, pecuniary-based incentives are scarcer in the public sector, partially because of insufficient funds for performance bonuses but also reflecting public aversion to paying public employees cash bonuses. Although recognition-based incentives might have quite different motivational mechanisms than monetary compensation, research suggests they may produce comparable, or in some cases superior, results. Recognition may bolster an agent’s intrinsic motivation or engender less negative reactions among coworkers than pecuniary awards (Frey & Benz, 2005). In Heinrich’s (2007) study of performance bonus awards in the Workforce Investment Act (WIA), she observed that performance bonuses paid to states tended to be low powered, in terms of both the proportion of performance-contingent funds to total program funding and the lack of relationship between WIA program employees’ performance and how bonus funds were expended by the states. Yet WIA program employees appeared to be highly motivated by the nonmonetary recognition and symbolic value associated with the publication of WIA’s high-performance bonus system results. In other words, despite the low or negligible monetary bonus potential, WIA employees still cared greatly about their measured performance outcomes and actively responded to WIA performance system incentives.

Job Design and Performance Measurement

Holmstrom and Milgrom (1991) explored job design as a means for controlling system incentives in their multitasking model in which the principal can divide responsibility for work tasks among agents and determine how performance will be compensated for each task. Particularly in the presence of agents with differing levels or types of motivation, the principal could use this information to group and assign different tasks to employees according to the complexity of tasks and how easily outcomes are observed and measured. Employees who are intrinsically motivated to supply high levels of effort and are more likely to act as stewards (identifying with the goals of the principal) could be assigned tasks that require a greater level of discretion and/or are more difficult to monitor and measure.⁷ For tasks in

⁷ Job assignment by preferences, of course, assumes that the principal can observe the agent’s preference type.

which performance is easily measured, the incentives provided for working hard and achieving high levels of performance should be stronger.

We consider the potential of job design to address an ongoing challenge in the performance-measurement systems of public programs—the selective enrollment of clients and/or limitations on their access to services—that are intended to increase measured performance, even in the absence of any value added by the program. In the U.S. Job Training Partnership Act (JTPA) program discussed above, job placement rates and wages at placement performance were straightforward to measure but also easy to “game” by enrolling individuals with strong work histories (e.g., someone recently laid off from a well-paying job). Partly in response to this problem, the JTPA program subsequently added separate performance measures for “hard-to-serve” groups, and its successor, the Workforce Investment Act (WIA) program, added an earnings change measure. As we discuss further below, the addition of these measures did not discourage strategic enrollments; rather, they primarily changed the groups who were affected by this behavior.

Since it is difficult to discourage this type of behavioral response, another possibility would be to separate work tasks, such as enrollment, training, and/or job placement activities performed by groups of employees. Prendergast’s (2007) research suggests that in this case, workers who are more intrinsically motivated should be assigned the tasks of training (i.e., increasing clients’ skills levels), as training effort is more difficult to evaluate and is more costly to exert. Low-powered incentives (or no incentives at all) should be tied to the training activity outputs. Alternatively, extrinsically motivated employees could undertake activities for which performance is easier to measure, such as job development activities and enrollment tasks that specify target goals for particular population subgroups; with their performance more readily observed, higher-powered incentives could be used to motivate these employees. Ideally, the measured performance of employees engaged in job placement activities would be adjusted for the skill levels achieved by clients in the training component of the program. Under WIA, the more intensive training services are currently delivered through individual training accounts (ITAs) or vouchers by registered providers, which, in principle, would facilitate a separation of tasks that could make a differential allocation of incentive power feasible. However, to realize the incentive system design described above, the registration of providers would have to be limited to those with primarily intrinsically motivated employees, which would effectively undermine the logic of introducing market-like incentives (through the use of vouchers) to encourage and reveal agent effort.

In the area of welfare services delivery, there has recently been more systematic study of the potential for separation of tasks to influence employee behavior and program outcomes. In her study of casework design in welfare-to-work programs, Hill (2006) explored empirically the implications for performance of the separation of harder-to-measure, core casework tasks (i.e., client needs assessments, the development of employability plans, arranging and coordinating services, and monitoring of clients’ progress) from other caseworker tasks that are more readily measured, such as the number of contacts with employers for job development or the number of jobs arranged. Her central hypothesis was that separating the stand-alone casework tasks (from the core tasks) into different jobs would lead to greater program effectiveness, as measured in terms of increasing client earnings and reducing AFDC receipt. Hill found support for her hypothesis “consistent with the incentive explanation from multitask principal agent theory—that is, staff are better directed toward organizational goals when unmeasurable and measurable tasks are grouped in separate jobs” (p. 278). The separation of casework tasks contributed positively to welfare-to-work program client earnings, although she did not find statistically significant associations with reductions in client receipt of AFDC benefits.

More often than not, though, it is simply not pragmatic to separate work tasks in complex production settings, and some administrative structures in the public sector

may only allow for joint or “team production.” In fact, there has been a shift toward the use of competitive monetary bonuses to motivate and reward performance at higher (aggregate) levels of government organization, such as schools or school districts in public education systems, job training centers under the Job Corps program, and states in the WIA and Temporary Assistance for Needy Families (TANF) programs. In these and other public programs, aggregate performance is calculated from client-level information, and performance bonuses are awarded based on rankings or other formulae for assessing the achievement of performance goals by the agency, state, or other aggregate unit. As Holmstrom (1982) explicated, the use of relative performance evaluation such as this will be valuable if one agent’s or group’s outcomes provide information about the state uncertainty of another, and only if agents face some common uncertainties.

For example, the postmaster general’s office supervises thousands of local post offices, each providing the same kinds of services but to different populations of clients. The performance of a local office’s peers (such as post offices in neighboring jurisdictions) will contain information about its effort, given that the performance of each is a function not only of the effort exerted, but also of external, random factors. Some of these external factors will be idiosyncratic, such as the weather, while others are shared, such as federal restrictions on postal rate charges. Relative performance evaluation—basing an agent’s performance on how it compares to his peers—works by differencing out the common external factors that affect all agents’ performance equally, which reduces some of the risk the agent faces and allows the principal to raise the incentive intensity, and subsequently, agents’ effort levels.⁸ It is one example of Holmstrom’s (1979) “informativeness principle”: The effectiveness of an agent’s performance measure can be improved by adjusting the agent’s performance by any secondary factor that contains information about the measurement error.

Another example of the informativeness principle is the use of past performance information to learn about the size of the random component and to reduce the noise in the measure of current effort, an approach that is currently being used in many school districts to measure the value added of an additional year of schooling on student achievement (Harris & Sass, 2006; Koedel & Betts, 2008; J. Rothstein, 2008).⁹ Roderick, Jacob, and Bryk (2004), for example, used a three-level model to estimate, at one level, within-student achievement over time; at a second level, student achievement growth as function of student characteristics and student participation in interventions; and at a third level, the estimated effects of any policy interventions on student achievement as a function of school characteristics and the environment. Thus, not only are measures of performance (i.e., student achievement) adjusted for past school performance, measurement error, and other factors over which educators have little control, but the models also attempt to explicitly estimate the contributions of policies and interventions that are directly controlled by educators and administrators to student outcomes.

⁸ While relative performance evaluation among agents within an organization may lead to a more precise estimate of agent effort, it may also promote counterproductive behavior in two ways. First, pitting agents against each other can reduce the incentive for agents to cooperate. In organizations where the gains from cooperation are great, the organization may not wish to use relative performance evaluation, even though it would produce a more precise estimate of agent effort. Second, pitting agents against one another increases the incentive for an agent to sabotage another agent’s performance. These behaviors would militate against the use of tournaments in some organizations.

⁹ Using an assessment period to create the performance standard creates an incentive for the agent to withhold effort. The agent will withhold effort because he realizes that exceptional performance in the assessment period will raise the bar and lower his compensation in subsequent periods, everything else equal. This is the so-called *ratchet* problem. One way to make use of the information contained in past performance while avoiding the ratchet problem is to rotate workers through tasks. By setting the standard for an agent based on the performance of the agent that occupied the position in the previous period, the perverse incentives that lead to the ratchet problem are eliminated.

Both of these examples of the application of the informativeness principle were also reflected in the U.S. Department of Labor's efforts under JTPA to establish expected performance levels using regression-based models. States were allowed to modify these models in calculating the performance of local service delivery areas, taking into account past performance and making adjustments for local area demographic characteristics and labor market conditions that might be correlated with performance outcomes. For example, performance expectations were lower for training centers located in relatively depressed labor market areas compared to those located in tighter labor markets. The aggregate unit for performance measurement was changed under WIA, however, to the state level; the states' aggregate performance relative to standards set for 17 different outcomes across three programs was calculated, with bonuses awarded based on states' performance levels, their relative performance (compared to other states), and performance improvements. In accord with Holmstrom's observations, the shift to a state-level performance measurement system has been problematic (see Heinrich, 2007). The local-level variation that was useful for understanding the role of uncertainties or uncontrollable events in performance outcomes is lost in the aggregation, and information gathered about performance is too far from the point of service delivery to be useful to program managers in understanding how their performance can be improved.

Dynamics Issues in Performance Measurement and Incentive Systems

As just described, the multitasking literature commonly employs a static perspective in modeling performance measurement design, which is applicable to both the public and private sectors. In practice, however, "real-world" incentive designers and policymakers typically begin with an imperfect understanding of agents' means for influencing measured performance and then learn over time about agents' behavior and modify the incentives agents face accordingly. For example, achievement test scores have long been used in public education by teachers, parents, and students to evaluate individual student achievement levels and grade progression, based on the supposition that test scores are accurate measures of student mastery of the subject tested, and that increases in test scores correlate with learning gains. More recently, however, under federal legislation intended to strengthen accountability for performance in education (No Child Left Behind Act of 2001), public schools are required to comprehensively test and report student performance, with higher stakes (primarily sanctions) attached to average student (and subgroup) performance levels.

Chicago Public Schools was one of the first school districts to appreciably increase the stakes for student performance on academic achievement tests, beginning nearly a decade before the pressures of the No Child Left Behind Act (NCLB) came to bear on all public schools. The early evidence from Chicago on the effects of these new systems of test-based accountability suggests some responses on the part of educators that were clearly unintended and jeopardized the goal of producing useful information for improving educational outcomes. In his study of Chicago Public Schools, Brian Jacob (2005) detected sharp gains in students' reading and math test scores that were not predicted by prior trajectories of achievement. His analysis of particular test question items showed that test score gains did not reflect broader skills improvement but rather increases in test-specific skills (related to curriculum alignment and test preparation), primarily in math. In addition, he found some evidence of shifting resources across subjects and increases in student special education and retention rates, particularly in low-achieving schools. And with Steven Levitt (Jacob & Levitt, 2003), he produced empirical evidence of outright cheating on the part of teachers, who systematically altered students' test forms to increase classroom test performance. The findings of these and related

studies call into question the use of student achievement test scores as a primary measure to guide efforts to improve the performance of public schools, although critics of current test-based accountability systems have yet to identify any alternative (and readily available) measures that are less noisy or prone to distortion.

Other examples suggest that incentive designers who discover dysfunctional behavior will often choose to replace or modify performance measures rather than eliminate incentive systems altogether. In the JTPA and WIA programs, policymakers added new performance measures (and discarded others) over time to reduce caseworkers' focus on immediate outcomes (job placement and wages at the time of program exit). JTPA's original set of performance measures also included cost measures—based on how much employees spent to produce a job placement—to encourage program managers to give weight to efficiency concerns in their decision making. However, eight years after their introduction, JTPA officials phased them out in response to research and experience showing that the cost standards were limiting the provision of longer-term or more intensive program services. In addition, federal officials noticed that local program administrators were failing to terminate (process the exit) of enrollees who, while no longer receiving services, were unemployed. By holding back poorly performing enrollees, even when these enrollees were no longer in contact with the program, training centers could boost their performance scores. Department of Labor officials subsequently closed this loophole by limiting the time an idle enrollee could stay in the program to 90 days (see Barnow & Smith, 2002; Courty & Marschke, 2007a; Courty, Kim, & Marschke, 2008).¹⁰

In effect, the experience of real-world incentive designers in both the public and private sectors offers evidence that the nature of a performance measure's distortions is frequently unknown to the incentive designer before implementation. In the public sector, the challenges may be greater because the incentive designer may be far removed from the frontline worker she is trying to manage and may have less detailed knowledge of the mechanics of the position and/or the technology of production. In the case of JTPA, Congress outlined in the Act the performance measures for JTPA managers, and staff at the U.S. Department of Labor refined them. It is unlikely that these parties had complete, first-hand knowledge of job training processes necessary to fully anticipate the distortions in the performance measures that they specified.

The Static Model Adapted to a Dynamic Context

Courty and Marschke (2003b, 2007b) adapt the standard principal-agent model to show how organizations manage performance measures when gaming is revealed over time. Their model assumes the principal has several imperfect performance measures to choose from, each capturing an aspect of the agent's productive effort, but each also is distorted to some extent. That is, the measures have their own distinct weaknesses that are unknown to the incentive designer *ex ante*. The model also assumes that the agent knows precisely how to exploit each weakness because of her day-to-day familiarity with the technology of production.

Courty and Marschke (2003b) assume multiple periods, and in any one period, the incentive designer is restricted to using only one performance measure. In the first period, the incentive designer chooses a performance measure at random (as the attributes of the measure are unknown *ex ante*) and a contract that specifies how

¹⁰ We thank a referee for pointing out that in WIA some states apparently allowed an interpretation of what constituted service receipt so that a brief phone call made to the enrollee every 90 days was sufficient to prevent an exit. In recent years, however, the Department of Labor clarified the definition of services so that (apparently) such phone calls and other kinds of non-training interactions with clients became insufficient to delay exit (Department of Labor, Employment and Training Administration Training and Employment Guidance Letter No. 15-03, December 10, 2003).

the agent's compensation will be determined by her performance on that measure. The model assumes the agent then chooses her effort to maximize her compensation in the current period net of effort costs (i.e., myopically). The incentive designer monitors the agent's actions and eventually learns the extent of the measure's distortion. If the performance measure proves to be less distorted than expected, the incentive designer increases the weight on the measure in the next period, and if the distortion is greater than expected, she replaces it with another measure. The model implies that major changes in performance measures are more likely to take place early in the implementation of performance measurement systems, before the incentive designer acquires knowledge of the feasible gaming strategies.

After the learning comes to an end, a dynamic learning model such as this delivers the same implications for the relations among incentive weights, alignment, risk aversion, and so on as does the static model. A new implication, however, is that the alignment between a performance measure and the true goal decreases as a performance measure is activated, or as it is more heavily rewarded, and increases as the measure is retired. This occurs because once the measure is activated, an agent will turn his focus on it as an objective and begin to explore *all* strategies for raising it—not just the ones that also improve organizational value.

For example, before public schools attached high stakes to achievement test outcomes, it indeed may have been the case that student test scores were strongly positively correlated with their educational achievement levels. Once educators came under strong pressures to increase student proficiency levels, increasing test scores rather than increasing student learning became the focal objective, and some educators began to look for quicker and easier ways to boost student performance. Some of these ways, such as teaching to the test, were effective in raising student test scores, but at the expense of students' learning in other academic areas. Thus, we expect that the correlation between student test scores and learning likely fell with the introduction of sanctions for inadequate yearly progress on these performance measures, generally reducing the value of test scores for assessing student learning and school performance.

The problem with using correlation (and other measures of association) to choose performance measures is that the correlation between an *inactivated* performance measure and value added does not pick up the distortion inherent in the performance measure. The change in the correlation structure before and after activation, as Courty and Marschke (2007b) show, reflects the amount of distortion in a performance measure and thus can be used to measure the distortion in a performance measure. They formally demonstrate that a measure is distorted if and only if the regression coefficient of value on measured performance decreases after the measure's activation;¹¹ then, using data from JTPA, they find that the estimated regression coefficients for some measures indeed fall with their introduction, suggesting the existence of performance measure distortion in JTPA. These findings corroborate their previous work (Courty & Marschke, 2004) that used the specific rules of the performance measurement system to demonstrate the existence of distortions in JTPA performance measures. In this work, the authors show that caseworkers use their discretion over the timing of enrollee exits to inflate their performance outcomes. In addition, the authors find evidence that this kind of gaming behavior resulted in lower program impacts.

This dynamic model shows that common methods used to select performance measures have important limitations. As Baker (2002) points out, researchers and practitioners frequently evaluate the usefulness of a performance measure based on how well the measure predicts the true organizational objective, using statistical correlation or other measures of association (see Ittner & Larcker, 1998; Banker, Potter, & Srinivasan, 2000; van Praag & Cools, 2001; Burghart et al., 2001; Heckman,

¹¹ They also show that a measure is distorted if its correlation with value decreases once it is applied.

Heinrich, & Smith, 2002). In practice, one has to know how much these correlations are capable of revealing about the agents' choices; if one assumes the linear formulation of noisy performance measures and value discussed earlier, a correlation of the two will pick up only the correlation in the noise (error) terms, which tells one nothing about the agent's choices. An additional implication of the dynamic model is that the association between a performance measure and organizational value is endogenous, casting further doubt on these methods for validating performance measures.

Learning in a Dynamic Context

In the dynamic model described above, the assumption of multiple periods introduces the possibility that agents will behave strategically. The model could be adapted to permit the agent to anticipate that the principal will replace performance measures once she learns that the agent is gaming. The main predictions of the model would still hold, however, as it would still be optimal for the agent to game.

The multitasking model could also be extended to permit learning by the agent. Koretz and Hamilton (2006) describe evidence that the difference between states' and school districts' performance on standardized tests and student learning appears to grow over time. Performance inflation is interpreted as evidence that agents (teachers) learn over time how to manipulate test scores. Thus, just as the static model could be reformulated to allow for the principal learning about how effective a performance measure is, it could be reformulated to allow the agent to learn how to control the performance measure. In such circumstances, little gaming would take place early on, but gaming would increase as the agent acquired experience and learned the measure-specific gaming technology. The amount of gaming that takes place in a given period, therefore, will depend on how distorted the measure was to start with, how long the agent has had to learn to game the measure, and the rate of learning. In other words, although a performance measure may be very successful when it is first used, its effectiveness may decline over time, and it may be optimal at some point to replace it.

In the Wisconsin Works (W-2) welfare-to-work program, for example, the state modified its performance-based contracts with private W-2 providers over successive (2-year) contract periods. In the first W-2 contract period (beginning in 1997), performance measures were primarily process oriented (e.g., the client-staff ratio and percent of clients with an employability plan). Provider profits in this period were tied primarily to low service costs, and welfare caseload declines that were already in progress made it easy for contractors to meet the performance requirements with little effort and expense; thus, W-2 providers realized unexpectedly large profits, from the state's perspective (Heinrich & Choi, 2007). The state responded in the second period by establishing five new outcome measures intended to increase contractors' focus on service *quality*. Profit levels were also restricted in this period, although W-2 providers continued to make sizable profits, partly because some performance targets were set too low to be meaningful, but also because contractors identified new ways to adjust client intake and activities to influence the new performance outcome levels. In the third contract period, the state added new quality measures and assigned different weights to them, with the "priority outcomes" performance measures accorded two-thirds of the weight in the bonus calculations. Again, the W-2 contractors appeared to grasp the new incentives and responded accordingly. Heinrich and Choi's analysis suggested that providers likely shifted their gaming efforts, such as changing participants' assignments to different "tiers" of W-2 participation that affected their inclusion in the performance calculations according to the weights placed on the different standards. In fact, the state decided to retire this and other measures in the fourth contract period and also rescinded the bonus monies that had been allocated to reward performance.

This example suggests that if an agent is learning along with the principal, the dynamics may get very complicated. And if performance measures degrade over time, then there will be a dynamic to measurement systems that will not necessarily stop. As discussed earlier, this type of dynamic also appears to be present in the selective enrollment of clients into the JTPA and WIA programs. In order to discourage enrollment practices strategically designed to increase performance (regardless of program value added), policymakers first developed separate performance measures for the “hard-to-serve” in the JTPA program. However, states and localities could choose to place more or less weight on these measures, and frontline workers learned over time how to select the most able from among the classes of hard-to-serve defined by the measures (e.g., the most employable among welfare recipients). The WIA program introduced universal access, yet access to more intensive training is still limited by a sequential enrollment process. Thus, WIA also added an earnings change performance measure intended to discourage the favorable treatment of clients with stronger preprogram labor force histories. After only four years, however, the U.S. Department of Labor recently concluded that it was encouraging a different type of strategic enrollment—discouraging service to workers with higher earnings at their last job in order to achieve a higher (earnings change) performance level. A U.S. GAO (2002, p. 14) report confirmed in interviews with WIA program administrators in 50 states the persistence of this type of behavioral response: “[T]he need to meet performance levels may be the driving factor in deciding who receives WIA-funded services at the local level.”

Setting and Adjusting Performance Standards

Another open dynamic issue concerns the growing number of performance measurement systems that have built requirements for “continuous performance improvement” or “stretch targets” into their incentive systems. A potential advantage of this incentive system feature is that it may encourage effort by agents to learn and improve in their use of the production technology, and this learning in turn increases the return at the margin of supplying effort. However, if agents’ efforts influence the rate at which performance improvements are expected, this could contribute to an unintended ratchet problem. In effect, if the agent learns or adjusts more quickly than the principal anticipates, this may lead to underprovision of effort.

In practice, government programs have been more likely to use a model with fixed increases in performance goals to set continuous performance improvement targets. In the first years of the WIA program, states established performance standards for three-year periods, using past performance data to set targets for the first year and building in predetermined expectations for performance improvements in the subsequent two years. Unfortunately, over the first three years of the program (July 2000–2003), an economic recession made it significantly more challenging for states to meet their performance targets. As performance standard adjustments for economic conditions are not standard practice in WIA, the unadjusted performance expectations were in effect higher than anticipated, which triggered a number of gaming responses intended to mitigate the unavoidable decline in performance relative to the standards (see Courty, Heinrich, & Marschke, 2005). The current U.S. Department of Labor strategic plan for fiscal years 2006–2011 has continued this approach, setting targets for six years, although a revised 2006 policy now allows local workforce development officials to renegotiate the performance targets in the case of “unanticipated circumstances” that make the goals unattainable.¹²

¹² See http://www.dol.gov/_sec/straplan/ and http://www.dwd.state.wi.us/dwdwia/PDF/policy_update0609.pdf (downloaded on May 26, 2009).

The use of “stretch” or benchmarked targets (i.e., typically based on past performance assessments) has been criticized for its lack of attention to or analysis of the process or capability for achieving continuous performance improvements. Castellano, Young, and Roehm (2004) identify the failure to understand variation in process or the inherent inability of a stable process to regularly achieve point-specific targets as a fatal flaw of performance measurement systems. In effect, if performance measurement is to produce accurate knowledge of the value added of government activities, then it is essential to understand both the relationship of government activity to performance outcomes and the factors that influence outcomes but cannot (or should not) be controlled by public managers. Estimates of performance are more likely to accurately (and usefully) reflect the contribution (or value added) of public managers and program activities to any changes in performance if performance expectations are adjusted for factors that are not controlled in production.

Yet in a recent review of the use of performance standard adjustments in government programs, Barnow and Heinrich (in press) show that their use is still relatively rare and concerns that adjustments are likely to be noisy, biased, and/or unreliable are prevalent. They also suggest, however, that such criticisms are likely overstated relative to the potential benefits of obtaining more useful information for improving program management and discouraging “gaming” responses (such as altering who is served and how to influence measured performance).¹³ The Job Corps program, for example, uses regression models to adjust performance standards for five of 11 performance measures that are weighted and aggregated to provide an overall report card grade (U.S. Department of Labor, 2001). The stated rationale is “by setting individualized goals that adjust for differences in key factors that are beyond the operator’s control, this helps to ‘level the playing field’ in assessing performance” (U.S. Department of Labor, 2001, Appendix 501, p. 6). Some of the adjustment factors include participant age, reading and numeracy scores, occupational group, and state and county economic characteristics. In practice, the Job Corps adjustment models frequently lead to large differences in performance standards across the centers. In 2007, the diploma/GED attainment rate standard ranged from 40.5 percent to 60.8 percent, and the average graduate hourly wage standard ranged from \$5.83 per hour to \$10.19 per hour. Job Corps staff report that center operators find the performance standards adjustments to be credible and fair in helping to account for varying circumstances across centers.

In fact, three different studies of the shift away from formal performance standard adjustment procedures under WIA concluded that the abandonment of the regression model adjustment process (used under JTPA) led to standards viewed as arbitrary, increased risk for program managers, greater incentives for creaming among potential participants, and other undesirable *post hoc* activities to improve measured performance (Heinrich, 2004; Social Policy Research Associates, 2004; Barnow & King, 2005). Barnow and Heinrich (in press) conclude that more experimentation with performance adjustments in public programs is needed, or we will continue to be limited in our ability to understand not only whether they have the potential to improve the accuracy of performance assessments, but also if they contribute to improved performance over time as public managers receive more useful feedback about their programs’ achievements (or failures) and what contributes to them.

Toward Better Systems for Measuring Performance Improvement

The fact that a majority of public sector performance measurement systems adopt a linear or “straight-line” approach to measuring performance improvement also

¹³ Courty, Kim, and Marschke (2008) show that the use of regression-based models in JTPA indeed seems to have influenced who in JTPA was served.

contributes importantly to the implementation problems described above. A straight-line approach typically defines a required (linear) rate of performance improvement from an initial score or target level and may also specify an ending value corresponding to a maximum performance level (e.g., the No Child Left Behind goal of 100% proficiency in reading and mathematics for public school students by 2014). Alternatively, in nonlinear performance incentive system designs, performance below a specified numerical standard receives no recognition or reward, and/or performance at or above the standard is rewarded with a fixed-size bonus.

A common flaw of straight-line performance improvement models for establishing performance expectations, explain Koretz and Hamilton (2006), is that they are “rarely based on empirical information or other evidence regarding what expectations would be realistic” (p. 540). Although this is also a concern in nonlinear incentive systems, the latter induce other distortions that are potentially more problematic. In their illustration of the principle of equal compensation in incentives, Holmstrom and Milgrom (1991) predicted that the implementation of nonlinear incentive contracts would cause agents to reallocate effort away from tasks or projects (or clients/subjects) for which greater effort is required to push the project or client over the threshold, or for which the project or client would exceed the threshold on their own (see also Holmstrom & Milgrom, 1987). In other words, agents would exert more effort in activities or with clients that were likely to influence performance close to the standard.

For example, in education systems with performance assessments based on specific standards for proficiency, we are more likely to observe a reallocation of effort to “bubble kids,” those on the cusp of reaching the standard (Pedulla et al., 2003; Koretz & Hamilton, 2006). Figlio and Rouse (2005) found that Florida schools that faced pressure to improve the fraction of students scoring above minimum levels on comprehensive assessment tests focused their attention on students in the lower portion of the achievement distribution and on the subjects/grades tested on the exams, at the expense of higher-achieving students and subjects not included in the test. In their study of school accountability systems in Chicago, Neal and Schanzenbach (2007) similarly concluded that a focus on a single proficiency standard would not bring about improvement in effort or instruction for *all* students, likely leaving out those both far below and far above the standard. The WIA example discussed earlier shows another form of this response. In the face of worsening economic conditions and the requirement to demonstrate performance improvements on an earnings change measure, case workers restricted access to services for dislocated workers for whom it would be more difficult to achieve the targeted increase in earnings (i.e., above their preprogram earnings levels).

The high-stakes elements of recent public sector performance incentive systems—such as the threat of school closure under the No Child Left Behind Act and the potential to win or lose funding or to be reorganized in public employment and training programs—have exacerbated these gaming responses and contributed to a growing interest in the use of value-added approaches in performance measurement. As discussed above, value-added models aim to measure productivity *growth* or improvement from one time point to another and the process by which any such growth or gains are produced, taking into account initial productivity or ability levels and adjusting for factors outside the control of administrators (e.g., student demographics in schools). Neal and Schanzenbach (2007) show that a value-added system with performance measured as the total net improvement in school achievement has the potential to equalize the marginal effort cost of increasing a given student’s test score, irrespective of the student’s initial ability. In effect, a value-added design eliminates (more than any alternative model) the incentive to strategically focus effort on particular subgroups of public program clients to improve measured performance.

Neal and Schanzenbach also acknowledge, however, that taking into account limits of or constraints on the technology of production, and in the case of public schools, students' initial distribution of ability, there may still be circumstances in which little effort would be exerted to help some student subgroups improve. Indeed, interest in implementing value-added performance measurement systems in public schools is expanding largely because they have the potential to produce useful information for policymakers in understanding how school/policy inputs and interventions relate to and can be better targeted to achieve net improvements in student outcomes (Roderick, Jacob, & Bryk, 2002, 2004; Thorn & Meyer, 2006). Described in terms of the value and performance models discussed above, it is a goal of the value-added approach to better understand and motivate the dimension of effort that is common to both performance measures and to value. Thorn and Meyer of the Value-Added Research Center at the Wisconsin Center for Education Research report how school district staff, with the support of researchers, are using value-added systems not only to measure teacher and classroom performance, but also to diagnose problems, track student achievement over time and assess the decay effects of interventions, and to remove incentives to teach narrowly to test outcomes or to focus on short-term gains. Of course, value-added performance measurement systems are still in the early stages of development and implementation in the public sector, and as in other incentive system designs, researchers are already identifying emerging measurement challenges and complex dynamics of which the implications will only become better known with time, experience, and careful study of their use (Harris and Sass, 2006; Lockwood et al., 2007; Koedel & Betts, 2008; J. Rothstein, 2008).

In fact, the development and expanding use of value-added systems for measuring performance in education appears to be encouraging exactly the type of "interactive dialogue" that Moynihan proposes to support more effective use of performance information and organizational learning. In February 2008, Vanderbilt University's National Center on Performance Incentives convened a conference that brought together federal, state, and local education policymakers, teachers and school support staff, and academics and school-based researchers to debate the general issues of incentives for improving educational performance, as well as the evidence in support of approaches such as value-added performance measurement systems and other strategies for assessing school performance. The conference organizers also invited participants who could illuminate lessons learned about performance measurement system design and implementation in other sectors, such as health care, municipal services, and employment and training (R. Rothstein, 2008). Two months later (April 2008) the Wisconsin Center for Education Research at the University of Wisconsin-Madison organized a conference to feature and discuss the latest research on developments in value-added performance measurement systems, which was likewise attended by both academic researchers and school district staff engaged in implementing and using value-added systems. As advocated by Moynihan (2008, p. 207), these efforts appear to be encouraging the use of value-added performance measurement systems as a "practical tool" for performance management and improvement, with "more realistic expectations" and "a clearer map of the difficulties and opportunities in implementation."

CONCLUSIONS

Drawing from examples of performance measurement systems in public and private organizations, we have sought to highlight key challenges in the design and implementation of performance measurement systems in the public sector. The complicated nature and technology of public programs compels performance measurement schemes to accommodate multiple and difficult-to-measure goals. One lesson we draw from the literature is that incentive schemes will be more effective if

they are not simply grafted onto the structure of an agency, regardless of its complexities or the heterogeneity of employees and work tasks. For example, an incentive system designer in a multitask environment, where some tasks are measurable and others are not, may be able to develop a more effective performance incentive scheme if care is taken to understand what motivates employees and to assign or reallocate tasks across workers accordingly. Assigning work so that one group of agents performs only measurable tasks and placing another group of intrinsically motivated workers in positions where performance is difficult to measure would exploit the motivating power of incentives for some workers and attenuate the moral hazard costs from the lack of incentives for the others. The usefulness of this strategy depends, of course, on the ability to identify intrinsically motivated workers and to facilitate a structural or functional separation of work tasks, which may be more or less feasible in some public sector settings.

Performance incentive system designers also need to appreciate and confront the evolutionary dynamic that appears to be an important feature of many performance measurement schemes. As illustrated, an incentive designer's understanding of the nature of a performance measure's distortions and employees' means for influencing performance is typically imperfect prior to implementation. It is only as performance measures are tried, evaluated, modified, and/or discarded that agents' responses become known. This type of performance measurement monitoring to assess a measure's effectiveness and distortion requires a considerable investment on the part of the principal, one that public sector program managers have neglected or underestimated in the past (Hefetz & Warner, 2004; Heinrich & Choi, 2007).

Of course, a complicating factor is that performance measures can be gamed. Agents (or employees), through their day-to-day experience with the technology of production, come to know the distinct weaknesses or distortions of performance measures and how they can exploit them. If it takes the agent time to learn how to game a new performance measure, an equilibrium solution may be a periodic revision of performance measures or a reassignment of agents. However, if both principals and agents are learning over time, the dynamic is likely to become more complex. And depending on the relative speeds of principal and agent learning and the extent to which performance measures degrade over time, this dynamic to measurement systems will not necessarily end. If the principal learns faster than the agent, the usefulness of a performance measure is more likely to increase, but if the agent learns faster how to manipulate a measure, its usefulness will decline and the measure may ultimately be discarded.

Furthermore, in a model with bilateral learning across multiple periods, not only can the form of the performance incentives have this never-ending dynamic, but the welfare path can also follow a very complicated course. For example, if agent learning about gaming is significant, one might realize an immediate improvement in social welfare with the implementation of performance incentives, followed by a drop in welfare as the agent learns how to game. If principal learning is important, an initial drop in welfare with the implementation of performance incentives may be followed by an improvement as the principal learns how to make adjustments that increase the efficiency of the system. And with a combination of principal learning and agent learning, numerous paths in welfare will be possible. This has important implications not only for the principal's decision about whether to implement an incentive system, but also for our ability to evaluate the effectiveness of incentive schemes. Whether or not one concludes that a particular incentive system is successful will likely depend on where in this dynamic cycle the program is being evaluated. We suggest that additional empirical and theoretical exploration of the dynamic aspects of performance measurement systems could be fruitful for generating additional insights and policy recommendations.

In practice, faced with these challenges, some incentive designers will “quit the game,” as in the recent case of the Temporary Assistance for Needy Families high-performance bonus system. At the same time, there are also promising examples of possible payoffs to persistence. The State of Wisconsin, in its fourth W-2 contract period, also appeared to give up on its performance incentive system, retracting its performance measures and funding allocations for bonuses. In the subsequent contract period, however, it made a comeback with the undertaking of a significant structural and management reorganization and new provisions for performance-based contracting. The state shifted to a four-year (2006–2009) contract to promote stability, more efficiently use administrative funds, and strengthen partnerships with other private service providers. In addition, it adopted new contract incentives that required larger W-2 contractors to “earn” 20 percent of their administrative and service contract funds through the achievement of specific outcomes; that is, reimbursement of costs was made conditional on these program outcomes being met. Perhaps most significant, the state broke its larger contracts and separately contracted out for W-2 case management, job development and placement, and Supplemental Security Income advocacy activities to allow for greater specialization, competition in service delivery, improved targeting of services to clients, and performance measures and incentives more closely aligned with service objectives. As with value-added performance measurement systems in education, only time will tell if the incentive designers have uncovered the right incentive scheme, although experiences to date in public employment and training programs suggest that the dynamic is unlikely to end in the current round.

Finally, the practice of measuring value or value added has been most fully developed in the context of performance accountability systems in education, but the idea of measuring value added (net of taxpayer contributions) is widely applicable in the public sector. The growing use of administrative data to track the type, timing, intensity, costs, and outcomes of public services delivery to individual citizens over time—whether in family or child welfare, employment and training, health care, transportation—has expanded the possibilities for attributing individual outcomes to receipt of services. Although, for most public programs, there is a long way to go before value-added or other nonexperimental analysis techniques that could be used with administrative data are formalized in performance management systems, recent studies suggest that with sufficient information on beneficiaries and variation in the type and duration of services received, there is potential to identify the effects of public services using administrative data alone. Heinrich, Mueser, and Troske (2008), for example, recently compared WIA program participants with matched comparisons of individuals who received employment services under Wagner–Peyser legislation or who received unemployment insurance benefits to evaluate the effects of WIA adult and dislocated worker programs in 12 states. Although such measures of value added (or program effects) are narrowly focused on quantitatively measured outcomes and are likely to be imperfect in their approximation of true value, we suggest that is worthwhile to explore their use and weigh their costs and benefits in performance management systems, with the perspective that they are just one of a number of possible tools for better understanding and managing public sector performance.

CAROLYN J. HEINRICH is Professor and Director at the La Follette School of Public Affairs, University of Wisconsin–Madison.

GERALD MARSCHKE is Associate Professor at the Department of Public Administration and Policy and Department of Economics, State University of New York at Albany, Faculty Research Fellow at NBER, and Research Fellow at IZA.

ACKNOWLEDGEMENTS

We thank Dan Koretz, George Baker, and other participants of the May 4, 2007, Eric M. Mindich Conference on Experimental Social Science, where this paper was first presented (at the Institute for Quantitative Social Sciences, Harvard University), for their valuable feedback. We also thank Doug Besharov, Maureen Pirog, and anonymous referees for their helpful comments and suggestions. Jerry Marschke gratefully acknowledges the hospitality of the Kennedy School of Government and the Labor and Worklife Program at Harvard Law School, where he was visiting while this paper was completed.

REFERENCES

- Akerlof, G., & Kranton, R. (2005). Identity and the economics of organizations. *Journal of Economic Perspectives*, 19, 9–32.
- Asch, B. J. (2005). The economic complexities of incentive reforms. In R. Klitgaard & P. C. Light (Eds.), *High-performance government: Structure, leadership, incentives* (pp. 309–342). Santa Monica, CA: RAND Corporation.
- Baker, G. P. (1992). Incentive contracts and performance measurement. *Journal of Political Economy*, 100, 598–614.
- Baker, G., Gibbons, R., & Murphy, K. J. (1994). Subjective performance measures in optimal incentive contracts. *Quarterly Journal of Economics*, 109, 1125–1156.
- Baker, G. P. (2002). Distortion and risk in optimal incentive contracts. *Journal of Human Resources*, 37, 728–751.
- Banker, R., Potter, G., & Srinivasan, D. (2000). An empirical investigation of an incentive plan that includes nonfinancial performance measures. *Accounting Review*, 75, 21–39.
- Barnow, B., & Smith, J. (2002). What does the evidence from employment and training programs reveal about the likely effects of ticket-to-work on service provider behavior? Unpublished manuscript.
- Barnow, B. A. (2000). Exploring the relationship between performance management and program impact: A case study of the JTPA. *Journal of Policy Analysis and Management*, 19, 118–141.
- Barnow, B. A., & Heinrich, C. J. (in press). One standard fits all? The pros and cons of performance standard adjustments.
- Barnow, B. S., & King, C. T. (2005). *The Workforce Investment Act in eight states*. Albany, NY: Nelson A. Rockefeller Institute of Government.
- Bertelli, A. M., & Lynn, L. E., Jr. (2006). *Madison's managers: Public administration and the Constitution*. Baltimore: Johns Hopkins University Press.
- Bevan, G., & Hood, C. (2006). What's measured is what matters: Targets and gaming in the English public health care system. *Public Administration*, 84, 517–538.
- Burgess, S., & Rato, M. (2003). The role of incentives in the public sector: Issues and evidence. *Oxford Review of Economic Policy*, 19, 285–300.
- Burghardt, J., Schochet, P., McConnell, S., Johnson, T., Gritz, M. R., Glazerman, S., Homrighausen, J., & Jackson, R. (2001). *Does Job Corps work? Summary of the National Job Corps Study (report 8140-530)*. Princeton, NJ: Mathematica Policy Research, Inc.
- Castellano, J. F., Young, S., & Roehm, H. A. (2004). The seven fatal flaws of performance measurement. *CPA Journal*, 74, 32–35.
- Courty, P., & Marschke, G. (2003a). Performance funding in federal agencies: A case study of a federal job training program. *Public Budgeting and Finance*, 23, 22–48.
- Courty, P., & Marschke, G. (2003b). Dynamics of performance-measurement systems. *Oxford Review of Economic Policy*, 19, 268–284.
- Courty, P., & Marschke, G. (2004). An empirical investigation of gaming responses to performance incentives. *Journal of Labor Economics* 22, 23–56.

- Courty, P., & Marschke, G. (2007a). Making government accountable: Lessons from a federal job training program. *Public Administration Review*, 67, 904–916.
- Courty, P., & Marschke, G. (2007b). A general test for distortions in performance measures. *Review of Economics and Statistics*, 90, 428–441.
- Courty, P., Heinrich, C. J., & Marschke, G. R. (2005). Setting the standard in performance measurement systems. *International Public Management Journal*, 8, 321–347.
- Courty, P., Kim, D. H., & Marschke, G. R. (2008). Curbing cream-skimming: Evidence on enrollment incentives. IZA Discussion Paper No. 3909.
- Davis, J. H., Donaldson, L., & Schoorman, F. D. (1997). Toward a stewardship theory of management. *Academy of Management Review*, 22, 20–47.
- Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125, 627–668.
- Deci, E. L., Koestner, R., & Ryan, R. M. (2001). Extrinsic rewards and intrinsic motivation in education: Reconsidered once again. *Review of Educational Research*, 71, 1–27.
- Dias, J. J., & Maynard-Moody, S. (2007). For-profit welfare: Contracts, conflicts, and the performance paradox. *Journal of Public Administration Research and Theory* 17, 189–211.
- Dickinson, K., West, R., Kogan, D., Drury, D., Franks, M., Schlichtmann, L., & Vencill, M. (1988). Evaluation of the effects of JTPA performance standards on clients, services, and costs. Research Report No. 88-16. Washington, DC: National Commission for Employment Policy.
- Dixit, A., (2002). Incentives and organizations in the public sector. *Journal of Human Resources*, 37, 696–727.
- Figlio, D. N., & Rouse, C. E. (2005). Do accountability and voucher threats improve low-performing schools? National Bureau of Economic Research Working Paper 11597.
- Francois, P., & Vlassopoulos, M. (2008). Pro-social motivation and the delivery of social services. Retrieved on October 26, 2008, from CESifo Economic Studies, <http://cesifo.oxfordjournals.org/cgi/content/full/ifn002v1>.
- Frederickson, D. G., & Frederickson, G. H. (2006). Measuring the performance of the hollow state. Washington, DC: Georgetown University Press.
- Frey, B. S., & Benz, M. (2005). Can private learn from public governance? *Economic Journal*, 115, F377–F396.
- Harris, D., & Sass, T. R. (2006). Value-added models and the measurement of teacher quality. Unpublished manuscript, Department of Economics, Florida State University, Tallahassee.
- Hatry, H. P. (1999). Performance measurement: Getting results. Washington, DC: Urban Institute Press.
- Heckman, J., Heinrich, C., & Smith, J. (2002). The performance of performance standards. *Journal of Human Resources*, 37, 778–811.
- Heckman, J., Smith, J., & Taber, C. (1996). What do bureaucrats do? The effects of performance standards and bureaucratic preferences on acceptance in the JTPA program. In G. Libecap (Ed.), *Advances in the study of entrepreneurship, innovation, and growth*, Vol. 7 (pp. 191–217). Greenwich, CT: JAI Press.
- Heckman, J. J., Heinrich, C. J., & Smith, J. A. (1997). Assessing the performance of performance standards in public bureaucracies. *American Economic Review*, 87, 389–395.
- Hefetz, A., & Warner, M. (2004). Privatization and its reverse: Explaining the dynamics of the government contracting process. *Journal of Public Administration Research and Theory* 14, 171–190.
- Heinrich, C. J. (1995). Public policy and methodological issues in the design and evaluation of employment and training programs at the service delivery area level. Doctoral dissertation, University of Chicago.
- Heinrich, C. J. (2004). Improving public-sector performance management: One step forward, two steps back? *Public Finance and Management*, 4, 317–351.
- Heinrich, C. J. (2007). False or fitting recognition? The use of high performance bonuses in motivating organizational achievements. *Journal of Policy Analysis and Management* 26, 281–304.

- Heinrich, C. J., & Choi, Y. (2007). Performance-based contracting in social welfare programs. *American Review of Public Administration* 37, 409–435.
- Heinrich, C. J., Mueser, P., & Troske, K. (2008). Workforce Investment Act non-experimental net impact evaluation. Final Report to U.S. Department of Labor.
- Hill, C. J. (2006). Casework job design and client outcomes in welfare-to-work offices. *Journal of Public Administration Research and Theory*, 16, 263–288.
- Holmstrom, B. (1979). Moral Hazard and Observability. *Bell Journal of Economics*, 10, 74–91.
- Holmstrom, B. (1982). Moral hazard in teams. *Bell Journal of Economics*, 13, 324–340.
- Holmstrom, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics and Organization*, 7, 24–52.
- Holmstrom, B., & Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55, 303–328.
- Houston, D. J. (2006). “Walking the walk” of public service motivation: Public employees and charitable gifts of time, blood, and money. *Journal of Public Administration Research and Theory*, 16, 67–86.
- Ittner, C. D., & Larcker, D. F. (1998). Are nonfinancial measures leading indicators of financial performance? An analysis of customer satisfaction. *Journal of Accounting Research*, 36, 1–35.
- Jacob, B. A. (2005). Accountability, incentives and behavior: Evidence from school reform in Chicago. *Journal of Public Economics* 89, 761–796.
- Jacob, B. A., & Lefgren, L. (2005). Principals as agents: Subjective performance measurement in education. Cambridge, MA: National Bureau of Economic Research Working Paper Number 11463.
- Jacob, B. A., & Levitt, S. D. (2003). Rotten apples: An investigation of the prevalence and predictors of teacher cheating. *Quarterly Journal of Economics*, 118, 843–877.
- Koedel, C., & Betts, J. (2008). Value-added to what? How a ceiling in the testing instrument influences value-added estimation. National Center on Performance Incentives Working Paper 2008-21, Vanderbilt University, Peabody College.
- Koretz, D., & Hamilton, L. S. (2006). Testing for accountability in K–12. In R. L. Brennan (Ed.), *Educational measurement*, 4th ed. (pp. 531–578). Westport, CT: American Council on Education/Praeger.
- Lockwood, J. R., McCaffrey, D. F., Hamilton, L. S., Stecher, B., Le, V.-L., & Martinez, J. F. (2007). The sensitivity of value-added teacher effect estimates to different mathematics achievement measures. *Journal of Educational Measurement*, 44, 47–67.
- Moynihan, D. P. (2008). *The dynamics of performance management: Constructing information and reform*. Washington DC: Georgetown University Press.
- Murdock, K. (2002). Intrinsic motivation and optimal incentive contracts. *RAND Journal of Economics*, 33, 650–671.
- Neal, D., & Schanzenbach, D. (2007). Left behind by design: Proficiency counts and test-based accountability. Unpublished manuscript, University of Chicago.
- Pedulla, J. J., Abrams, L. M., Madaus, G. F., Russell, M. K., Ramos, M. A., & Miao, J. (2003). *Perceived effects of state-mandated testing programs on teaching and learning: Findings from a national survey of teachers*. Boston: National Board on Educational Testing and Public Policy.
- Perry, J. L., & Porter, L. (1982). Factors affecting the context for motivation in public organizations. *Academy of Management Review*, 7, 89–98.
- Perry, J. L., & Wise, L. R. (1990). The motivational bases of public service. *Public Administration Review*, 50, 367–373.
- Prendergast, C. (2007). The motivation and bias of bureaucrats. *American Economic Review* 97, 180–196.
- Radin, B. (2006). *Challenging the performance movement*. Washington, DC: Georgetown University Press.

- Rainey, H. G. (1983). Public agencies and private firms: Incentive structures, goals, and individual roles. *Administration and Society*, 15, 207–242.
- Rainey, H. G. (2006). Reform trends at the federal level with implications for the states: The pursuit of flexibility and the human capital movement. In J. E. Kellough & L. G. Nigro, (Eds.), *Civil service reform in the states: Personnel policies and politics at the sub-national level* (pp. 33–58) Albany: State University of New York Press.
- Roderick, M., Jacob, B. A., & Bryk, A. S. (2002). The impact of high-stakes testing in Chicago on student achievement in the promotional gate grades. *Educational Evaluation and Policy Analysis*, 24, 333–357.
- Roderick, M., Jacob, B. A., & Bryk, A. S. (2004). Summer in the city: Achievement gains in Chicago's summer bridge program. In G. D. Borman & M. Boulay (Eds.), *Summer learning: Research, policies, and programs* (pp. 73–102). Mahwah, NJ: Lawrence Erlbaum Associates.
- Rothstein, J. (2008). Teacher quality in educational production: Tracking, decay, and student achievement. Unpublished Manuscript, Princeton University.
- Rothstein, R. (2008). Holding accountability to account: How scholarship and experience in other fields inform exploration of performance incentives in education. National Center on Performance Incentives Working paper 2008-04, Vanderbilt University, Peabody College.
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York: D. Appleton-Century.
- Social Policy Research Associates. (2004). *The Workforce Investment Act after five years: Results from the national evaluation of the implementation of WIA*. Oakland, CA: Social Policy Research Associates.
- Stone, D. (1997). *Policy paradox: The art of policy decision making*. New York: W. W. Norton.
- Taylor, F. (1911). Principles and methods of scientific management. *Journal of Accountancy*, 12, 117–124.
- Thorn, C. A., & Meyer, R. H. (2006). Longitudinal data systems to support data-informed decision making: A tri-state partnership between Michigan, Minnesota and Wisconsin. Report of the Wisconsin Center for Education Research.
- U.S. Department of Labor, Office of Job Corps. (2001). *Policy and Requirements Handbook: Job Corps*. Washington, DC: U.S. Department of Labor, Office of Job Corps.
- U.S. Government Accounting Office. (2002). Improvements needed in performance measures to provide a more accurate picture of WIA's effectiveness. GAO Report No. 02-275.
- van Praag, M., & Cools, K. (2001). Performance measure selection: Noise reduction and goal alignment. Unpublished manuscript, University of Amsterdam.
- Van Slyke, D. M. (2007). Agents or stewards: Using theory to understand the government-nonprofit social service contracting relationship. *Journal of Public Administration Research and Theory*, 17, 157–187.
- Walker, R. M., & Boyne, G. A. (2006). Public management reform and organizational performance: An empirical assessment of the U.K. Labour government's public service improvement strategy. *Journal of Policy Analysis and Management*, 25, 371–393.