# Short-run Incentives and Myopic Behavior: Evidence from State-Owned Enterprises in China

Peter Zeitz[1], UCLA, zeitzp@gmail.com, November 23, 2009
Job Market Paper

Abstract

How do links between pay and performance affect firm productivity? In 1978, Chinese industrial planners carried out major reforms of the compensation system in state-owned enterprises (SOEs), introducing bonuses which related pay to the fulfillment of short-run objectives. Previous studies have argued that strengthening incentives increased effort levels and led to large increases in productivity. However, incentives could have encouraged enterprises to misallocate effort by placing excessive emphasis on short-run targets and neglecting long-range activities. To test whether incentives led to myopic behavior, I collect a unique panel of compensation, employment, and output statistics and use these data to estimate the short-run and long-run effects of incentive use on labor productivity during the early period of SOE reform, 1976 to 1988. The results indicate that productivity gains were due to effort misallocation, rather than improvements in management. In the data, incentives were associated with a small increase in labor productivity in the short-run, and a large decrease in labor productivity in the long-run.

## Introduction

In 1978, Chinese industrial planners reformed the compensation system in state-owned enterprises (SOEs), introducing bonuses which linked employee wages to the fulfillment of performance targets. Beginning in 1979 and continuing through the 1980s, incentive reforms were deepened by linking the total sum available for bonuses and wage increases to enterprise performance. Previous researchers have argued that these links increased effort levels in SOEs and improved enterprise productivity, but there are reasons to be skeptical of this interpretation. Since bonus distributions were based on observations of enterprise performance, the causal relationship between bonuses and productivity is difficult to disentangle. Past studies have not addressed this issue sufficiently. More importantly, tying wages to short-run outcomes does not necessarily strengthen incentives to improve long-run productivity. Past efforts in long-run activities such as maintenance, research, and planning are important inputs in firms' current period production. Tying pay to short-run performance could encourage enterprises to neglect these activities in order to devote more effort to meeting current period targets. If this is the case, increases in short-run output associated with bonus use might have involved the sacrifice of future performance.

[1]Naomi Lamoreaux provided invaluable guidance throughout the project. Particularly useful comments and suggestions were received from Daniel Ackerberg, Leah Boustan, Dora Costa, Harold Demsetz, Steven Nafziger, John Riley, Connan Snider, and R. Bin Wong. The work was supported by an NSF doctoral dissertation improvement grant, an IIE Fulbright fellowship, and a UC Pacific Rim Grant.

At the onset of the reform period, state-owned enterprises stood at the center of China's planned economy, producing 76 percent of gross industrial output and accounting for 57 percent of industrial employment in 1978. As part of the 1980s reform process, SOEs were granted the right to market any output they produced in excess of quotas set under the planning system. Simultaneously, the state allowed SOEs to distribute a portion of firm profits to managers and workers as bonus pay. Though SOE privatization and the growth of other sectors have reduced their relative importance, SOEs remain significant to the Chinese economy. Even as late as 2004, 38 percent of China's industrial output was produced in state-owned firms or corporations in which government entities held a majority stake (Li and Putterman 2008). As a result, the extent to which the provision of incentives for workers and managers can improve SOE productivity remains an important question.

Previous research has found that the introduction of bonus pay during the 1980s resulted in a large improvement in SOE productivity (Groves et al. 1994, Li 1997, Yao 1997). However, none of these studies have investigated whether productivity increases persisted over time. Unlike these authors, I distinguish between the short- and long-run effects of bonus use on productivity. I measure long-run effects by allowing the use of bonuses over the past 10 years to influence current labor productivity. Incorporating past bonus use as a productivity determinant requires information spanning a long period. I collect a unique panel of compensation, employment, and output statistics from 977 SOEs in the Chinese iron and steel industry, operating between 1949 and 1988, and use these data to evaluate the effect of bonuses on labor productivity from 1976 to 1988. A comparison of short- and long-run effects allows for a reevaluation of the supposed benefits associated with the introduction of bonuses. I find that the introduction of bonuses led to a small increase in average labor productivity levels from 1978 to 1982, but a large decrease in labor productivity from 1984 to 1988.

Whether short-run incentives improve productivity depends on the activities workers perform. Workers subject to intense incentives often focus on rewarded activities to the exclusion of all else (Kerr 1975). For workers responsible for a single task, intense focus is generally beneficial. However, incentives may be less effective at motivating employees to take a holistic interest in their firm's success, particularly where they have multiple responsibilities. This problem was first formalized in Holmstrom and Milgrom (1991) who argued that the structure of conventional agency models, which assume that agents perform one task only, rather than 'multitask', had led economists to ignore important issues. In the Holmstrom and Milgrom model, employees face a decision over how much effort to supply and how to allocate effort across multiple productive activities. The employer benefits from all of these activities, but can only observe performance in a subset of them. Basing wage payments on these observations induces workers to withdraw effort from areas where performance is unobserved.

Some empirical studies have demonstrated that incentives can raise effort levels, while others have shown that incentives encourage suboptimal work strategies. Lazear (2000) analyzes a windshield installation firm where labor productivity improved dramatically as a result of a shift from time to piece wages. Similarly, Shearer (2004) finds that tree planters are more productive under piece rate pay than time wages. Other authors have focused on more complex activities, where agents have more opportunities to engage in strategic behavior. In the Soviet Union, managers earned bonuses from 30 to 100 percent of their base salary for meeting monthly production quotas, and were frequently dismissed for failure. Berliner (1956) argues that incentives created a vicious cycle of ‘storming’ to meet targets at the end of each quota period, and led to poor maintenance practices. Oyer (1998) documents similar behaviors among US managers and sales workers who discount products at the end of quota periods, which allows them to fulfill revenue targets. Asch (1990) studies the behavior of navy recruiters subject to periodic performance assessments, and finds that recruiting success peaks at assessment time and falls off subsequently.

While these studies do demonstrate that incentives have both positive and negative effects, they provide little information about their relative magnitude. This information is crucial to understanding why some firms choose to use time wages to motivate workers, while others offer incentive payments. One way of addressing this issue is through comparison of the effect of incentives on performance in the short- and long-run. Multitasking theory proposes that rewards based on current performance induce agents to shift effort away from activities which improve future performance. In large part, this argument is based on the intrinsic difficulty in measuring performance in long-run activities such as maintenance, research, and planning. When a considerable lag intervenes between when an activity occurs and its effects become evident, it can be difficult to identify the responsible individual and offer him an appropriate reward. This could lead principals to offer time wages and tolerate shirking when they hire agents to engage in long-run activities.

The causal effects of shifts in incentive practices are difficult to identify empirically because compensation practices are endogenous to firm and workforce characteristics. Private firms probably adopt performance-based pay because they anticipate better results from this policy than firms which do not adopt performance-based pay. This type of selection is difficult to control for because it is based on unobservable characteristics. To resolve this issue, it is necessary to identify external circumstances which affect compensation practices independently of firm level characteristics, and have a varied impact across firms. In market settings, this is typically impossible because compensation decisions are generally at the firm’s discretion.

The Chinese case is exceptional because the decision to use performance-based pay was imposed by the government rather than chosen by firms. In 1978, provincial ministries carried out pilot schemes which limited the use of incentives to experimental groups of firms (Liaoning Wage Policy 1978). Subsequently, beginning with a pilot program in Sichuan in 1979, the government granted firms rights to adjust incentive structures, but their autonomy was greatly limited by regulations which varied across provinces (Light Industry Research Group 1980, Koyima 1987). In both pilot programs, provincial planners selected relatively high-performing firms for participation to ensure that the programs would appear successful (Shirk 1993). Later, in the early 1980s, pilot schemes were extended to the general population of enterprises, and by the mid-1980s adoption of incentive pay was essentially universal. Importantly, bonus policies varied across provinces and industries, and over time, meaning that the intensity of the incentives used was determined primarily by the firm’s regulatory environment and not by the anticipated effectiveness of incentives in individual enterprises. This makes the Chinese SOE system an exceptionally good environment to test for causal effects of incentives on firm performance.

To test whether short-run incentives induced myopic behavior, I estimate the effect of current and past use of bonuses on current labor productivity. I use the share of wages distributed as performance bonuses as a measure of incentive intensity. A major problem here is that the bonus share is influenced by whether firms meet their performance targets. This makes the bonus share an endogenous variable; correlation between current performance and the bonus share will include a component due to reverse causation. I control for endogeneity using three separate methods, all of which yield similar results. In the first method, I make assumptions about bonus setting policy and the information observed by planners which allow the causal effect of bonuses to be separated from the effect due to reverse causation. Using an OLS regression based on these assumptions, I estimate the causal effect of current and past bonus use on labor productivity. In the second method, I exploit planners' selection of firms for the experimental introduction of incentive pay in 1978 as a source of exogenous variation. I compute a matching estimator which quantifies the effect of participation on short- and long-run growth in labor productivity. In the third method, I incorporate information on firm wage levels in OLS regressions to compare the effect of increases in bonus pay to that of increases in time wages. Like changes in bonus pay, wage levels are regulated by planners on the basis of firm performance, and thus serve as a proxy for unobserved performance determinants. This greatly reduces the potential for reverse causation to contaminate the estimated effect of bonus use.

Estimates from each technique indicate that firms which relied on bonuses had superior performance in the current period, but inferior performance in the long-run. Large and negative long-run effects imply that bonus use had an adverse impact on overall performance in the Chinese iron and steel industry. To assess this effect, I use regression estimates to calculate the implied impact of actual bonus use on labor productivity in 1988 relative to the counterfactual that bonus use was avoided entirely. These calculations indicate that the use of bonuses decreased 1988 labor productivity by between 0.08 and 0.25 log points, with the magnitude of the estimate depending on which regression specification is used to compute the counterfactual. Matching estimates of the treatment effect of adopting incentives in 1978 suggest even larger negative effects. Firms which adopted incentives in 1978 (the treated group) had bonus shares during the 1980s which averaged only 18 percent greater than those of firms which did not adopt incentives at this time (the untreated group). Despite the modest size of the treatment, estimates indicate that the performance of the treated group was significantly inferior to that of untreated firms. Between 1978 and 1988, labor productivity growth among treated firms was 0.17 log points lower than growth in the untreated group.

I divide my discussion into six sections. First, I present a simple example of a multitasking agency relationship in which excessive rewards for current performance distort decision-making and lead to decreases in long-run output. I relate this prediction to the behavior of a Chinese iron and steel firm which appears to have responded to incentives as the model predicts. This raises the question of why such rewards may be put in place even when they lead to inefficient outcomes. In large organizations, owners often cede control to a manager, and evaluate him based on the firm's current performance. In this situation, the manager can improve his apparent performance by establishing inefficient incentive practices. Next, I discuss governance problems in the Chinese SOE system which made it difficult for the central government, the principal in this setting, to control enterprise-level policies and encouraged lower level government agencies to establish inefficient incentive practices. The third section discusses methodological problems in previous studies of bonus use in Chinese SOEs, in particular their lack of adequate controls for reverse causation and their failure to distinguish short- and long-run effects. One of the factors hindering previous studies has been the lack of a set of panel data with an adequate time dimension. In the fourth section, I describe the collection of a novel large-scale dataset which makes it possible to study incentives in a dynamic context. The fifth section places changes in incentive use, and other policies adopted during the 1980s, in the context of policy changes occurring in prior decades. I argue that, to a significant extent, improvements in SOE performance during the 1980s reflected improvements in social stability and the resumption of earlier policies dating from the 1950s, rather than the effects of incentives. The sixth section presents statistical results indicating that performance-based compensation had a negative effect on long-run labor productivity. Finally, I conclude the paper by reviewing my findings and their implications for the study of incentive use in firms.

## A Simple Multitasking Example

Before considering the specifics of the incentive policy in China, I construct a very simple multitasking problem which illustrates why production incentives can have negative effects on future performance. The model is in the same spirit as Holmstrom and Milgrom (1991) in that agents respond to incentives by misallocating effort, but differs in that I assume that agents are intrinsically motivated. Other papers considering intrinsic motivation and incentive use include Benabou and Tirole (2003), Prendergast (2008), and Besley and Ghatak (2005). All of these papers analyze settings where intrinsic motivation and incentives are substitutes, whereas I analyze a case where they are complementary. In my model, highly-motivated agents are less prone to misallocate effort when tempted with an incentive, and thus the principal can offer them much larger incentives. Previous researchers have argued that incentives had a powerful positive effect in China because initial motivation levels were quite low (Groves et al. 1994). The model considers a case where strong incentives have negative effects in the presence of weak motivation, but positive effects when offered to motivated agents. The model follows:

A principal hires an agent to undertake production and investment for a finite period. The principal's goal is to induce the agent to maximize the sum of production and investment. Output in production is contractible, whereas output in investment is not observable and thus not contractible. The agent is intrinsically motivated, meaning that he receives utility which is a fraction $m \in [0,1]$ of the sum of production and investment. As an additional motivator, the principal can offer the agent a bonus, $b \in [0,1]$, which is a share of production output. Basing his decision on the size of the bonus offered, the agent allocates an effort level, $e_q \in R_+$, to production, and an effort level, $e_i \in R_+$, to investment. Production and investment output, $Q \in R$ and $I \in R$, respectively, are logarithmic functions of effort which are allowed to take negative values. Negative investment can be interpreted as the transformation of the firm's existing capital stock into production output. Total output, denoted Y, is a concave function of the agent's effort levels as shown in equation (1).

$$\mathrm{Y}(\mathrm{e_q}, \mathrm{e_i}) = \mathrm{Q}(\mathrm{e_q}) + \mathrm{I}(\mathrm{e_i}) = \ln \mathrm{e_q} + \ln \mathrm{e_i} \qquad (1)$$

The agent incurs an effort cost C which is a convex function of the agent's effort levels as shown in equation (2).

$$C(\mathrm{e_q}, \mathrm{e_i}) = \begin{cases} 0 & if \quad \mathrm{e_q} + \mathrm{e_i} < 2 \\ \frac{1}{2}(\mathrm{e_q} + \mathrm{e_i} - 2)^2 & otherwise \end{cases} \qquad (2)$$

The agent's problem is to choose effort levels which maximize his utility, as shown in equation (3).

$$\max_{\mathrm{e_q}, \mathrm{e_i}} \{(\mathrm{m} + \mathrm{b})\mathrm{Q}(\mathrm{e_q}) + \mathrm{mI}(\mathrm{e_i}) - C(\mathrm{e_q}, \mathrm{e_i})\} \qquad (3)$$

Since the agent's utility is a concave function of his effort level, his optimal effort choices are given by the first order conditions of equation (3). To aid in interpretation, I list the agent's optimal effort level for each activity, together with his aggregate effort level, e, in equation (4).

$$\mathrm{e}(\mathrm{m}, \mathrm{b}) = \mathrm{e_q}(\mathrm{m}, \mathrm{b}) + \mathrm{e_i}(\mathrm{m}, \mathrm{b}) = \left(1 + \sqrt{2m + b + 1}\right)$$

$$\mathrm{e_q}(m, b) = \frac{m+b}{2m+b} e(m, b) \ : \mathrm{e_i}(m, b) = \frac{m}{2m+b} e(m, b) \qquad (4)$$

The conditions in equation (4) indicate that the principal faces a trade-off. Increasing the bonus level will increase the agent's overall effort level, but also induce him to reduce investment effort. The reduction in investment is particularly drastic when the agent is poorly motivated, and thus places a low weight on investment output. The principal's objective is to choose a value of $b \in [0,1]$ which maximizes total output, as shown in equation (5).

$$\max_{\mathrm{b}} \{\mathrm{Y}(\mathrm{e_q}(\mathrm{m}, \mathrm{b}), \mathrm{e_i}(\mathrm{m}, \mathrm{b}))\} = \max_{\mathrm{b}} \left\{\ln\left(\frac{\mathrm{e_q}(m,b)\mathrm{e_i}(m,b)}{(e(m,b))}\right) + \ln(e(m, b))\right\} \qquad (5)$$

In equation (5), the first term, $\ln\left(\frac{\mathrm{e_q}(m,b)\mathrm{e_i}(m,b)}{(e(m,b))}\right)$, which is decreasing in b, measures loss in aggregate output from misallocation. The second term, $\ln(\mathrm{e}(\mathrm{m}, \mathrm{b}))$, measures gains in aggregate output from increases in the agent's overall effort level, $e(m, b)$. The first order condition for maximizing output implies that the principal chooses a bonus level of $b = m^2$.[2] This indicates that the optimal strength of incentives is increasing in the agent's motivation level, m. Since motivated agents internalize more of the costs of misallocation, principals can offer them more intense incentives. This is relevant to the context of Chinese SOEs because SOE managers and workers are often regarded as unmotivated and prone to shirking. The multitasking model presented here indicates that offering intense incentives to poorly motivated agents could lead to shirking in long-run activities, and thus that bonus provision could have had negative effects on SOEs' overall performance.

---

[2]The principal's objective function, given in equation (5), is not concave on $b \in [0,1]$, but is quasi-concave on this interval. To verify quasi-concavity and thus that $b = m^2$ is the unique maximum, note that

$$\frac{\mathrm{dY}}{\mathrm{db}} = \frac{-2(b-m^2) - b\left(\sqrt{(m+1)^2 + (b-m^2)} - (m+1)\right)}{(b+m)(b+2m)(\sqrt{b+2m+1})(1+\sqrt{b+2m+1})}$$

is strictly positive for all $b \in [0, m^2)$, equal to zero for $b = m^2$, and strictly negative for all $b \in (m^2, 1]$.

To illustrate the applicability of this model to the Chinese bonus system, it is worthwhile to discuss the experience of the Handan Steel Company as described in collected papers from a steel industry conference (The Reform of Small and Medium Iron and Steel Enterprises 1987). The incentive system at Handan was similar to that at most SOEs. From the early 1980s, Handan allocated a fixed fraction of its after-tax profits to a bonus fund, with the fraction set by province-level guidelines. Handan's managers could design some incentives for their own employees, but distribution of these rewards was always contingent upon the firm meeting monthly profit and sales targets set by provincial authorities. In the mid-1980s, Handan's managers were struggling to meet these targets and faced a dilemma. Their existing investment plan called for the completion of a steel rolling facility using internally produced materials. Managers anticipated that continued investment in the facility would make it impossible for the firm to meet profit and sales targets, and thus distribute bonuses. To avoid this, managers decided to put the project on hiatus and sell off construction materials. The revenue generated, either through the sale or through the reallocation of labor, allowed the firm to fulfill production targets and distribute bonuses. This outcome is consistent with the basic prediction of the model; Handan withdrew resources from investment to meet production targets. Implicitly, such a transfer should boost current output, but detract from output in the future.

SOE Governance in China

Multitasking models show that strong incentives can be inefficient, but do not explain why a firm might choose to implement them anyway. Models often assume that firm owners are sufficiently informed to control compensation practices directly. In practice, owners delegate control rights to a corporate board to represent their interests. Board members often have personal connections to managers, which can limit their willingness to aggressively pursue the interests of firm owners. Under weak oversight, managers may be able to negotiate a compensation contract which rewards self-dealing. Since the influence of long-term activities on firm value is intrinsically difficult to assess, compensation contracts tend to reward increases in short-run earnings more heavily than changes in firm value. The danger is that strong incentives to increase earnings will encourage managers to undertake activities that produce short-run results, but do not create value for owners.

In the Chinese SOE system, ultimate ownership resides with the central government, and managerial supervision is performed by province-level state agencies. Since future performance is difficult to assess, bureaucrats employed in supervisory agencies are evaluated for promotion on the basis of the current performance of the enterprises they supervise. Firm managers themselves are often selected from within the province-level bureaucracy. Given their tight links to management and their lack of an ownership interest, bureaucrats may choose to implement short-run incentives which encourage managers to engineer short-run improvements in firm performance, either by making short-run decisions or manipulating financial data. For the central government, it may be very difficult to distinguish a short-run increase in a firm's apparent performance from a lasting contribution to firm value. As a result, the central government is ill-equipped to assess the effects of incentive policies.

Besides regulating the level of incentive use, provincial agencies are also responsible for ensuring that enterprises meet performance targets. Though promotion opportunities encourage agency bureaucrats to perform this function, they may be less scrupulous about monitoring enterprise performance in investment-related areas. In order to fulfill short-term performance targets, managers often sacrifice the future of the enterprise by engaging in asset-value-reducing activities (Qian 1996). Bureaucrats may be unwilling to reveal these activities to the central government because their career prospects are linked to short-run enterprise performance, and because they have strong personal connections to enterprise managers.

Existing Studies of the 1980s Incentive Reform: The Endogeneity Problem

Existing studies have found that incentives had a positive effect on performance, but these studies suffer from some methodological problems. Econometric studies using the 1980-1989 Chinese Academy of Social Sciences (CASS) survey of 769 SOEs in four provinces, including those of Groves et al. (1994), Li (1997), and Yao (1997), have found that the introduction of bonuses improved enterprise productivity. The CASS data are based on responses to a survey sent out by the State Reform Commission to enterprises and includes firms from most industries. One concern with the use of the CASS data is that participating firms were not selected randomly; participants were larger than average and had faster labor productivity growth than the national SOE average. The reform effects which researchers are interested in are those in the general population of SOEs, but these could potentially differ significantly from those in samples selected by Chinese government agencies.

Studies using the CASS data estimate production functions incorporating measures of bonus use and other institutional controls as productivity determinants. Since planners essentially distribute bonuses across firms on the basis of performance, the potential for reverse causation in these regressions is serious; the failure to observe all factors affecting productivity could easily lead to a positive coefficient on bonuses even when they are irrelevant or harmful to performance. Solving this problem requires identifying a valid instrument for bonus use, but the instruments proposed in most studies are not convincing. Groves et al. (1994) and Yao (1997), for example, use a one-period lag of bonus payments as an instrument, but since a firm's recent performance—as reflected in recent bonus payouts—is positively correlated with current performance, this strategy is likely to bias estimates of the effects of bonus use upwards. Li (1997) argues that changes in planner-determined output quotas are set through a political process and not on the basis of firm productivity, and uses these allocations as instruments. However, this argument seems dubious because both managers and planners observe performance-related information and must factor this in when negotiating output quotas.

A second, potentially even more important, problem with these studies is that they focus only on short-run outcomes. The focus of past studies on the short run could be due to the limited time coverage of the CASS dataset, which includes only ten years of information. In a short panel, incorporating a firm's historical use of incentives as a performance determinant greatly reduces the sample size, and thus the power to test hypotheses. Omitting historical data on incentive use, however, may yield misleading results. To resolve this issue, I collect a very long panel containing up to 40 years of information, which allows me to condition on each firm's historical compensation practices without significantly reducing the sample size.

## The Dataset

The data I collect come from Chinese sources which have never been used before by Western scholars. For ease of reading, I translate their titles into English in the text; Chinese titles are given in the bibliography. The principal source of data is a three-volume set of plant-level wage statistics for the iron and steel industry, *Chinese Smelting Industry Labor and Wage Statistics, 1949-1988, Vols. 1-3* (1990), hereafter CSILWS, which was compiled from firm records by the Ministry of Metallurgy for internal government studies of the effectiveness of performance incentives. The sample includes almost every SOE which operated in the iron and steel industry in 1988, but excludes a very small number of firms which exited the industry prior to this year. Due to concern about the reliability of output measures from design and construction firms, I focus on 977 mining and manufacturing firms which produce standardized commodities, excluding 123 units engaged in planning, construction, research, and education. The range of commodities produced includes mineral ores such as iron, manganese, and fluorite, intermediate products such as coke, pig iron, and steel ingots, and final goods such as finished steel and smelting machinery.

The CSILWS data contain observations of gross output per worker, measured in constant 1980 prices, together with physical data on labor productivity measured in the tonnage of the firm's three principal outputs. Since most firms produced multiple outputs, I use gross value of output per worker as a performance measure. In theory, value-added is a better measure of output than gross value. However, cost accounting practices in China are poor, and SOEs are widely believed to manipulate their cost accounts by misreporting the price and usage of material inputs. The gross value of output is based on physical quantities, and is harder to manipulate than value-added information which relies on firm-level cost accounting. I control for variation in the value-added content across products through the use of fixed effects, product variety-specific time dummies, and (in some specifications) firm-specific time trends. Product variety-specific time trends also capture price changes which may affect firms differently depending on which product they produce. Wage data indicate total compensation, as well as the amounts paid on a time, piece, and bonus basis. I analyze the incentive effects of both piece rates and bonuses, but focus primarily on bonuses since they were the predominant form of performance-based pay during the reform period. Employment data includes total employment, employee sex, occupation (management/supervisory, technical, shop floor worker, trainee, service), and contractual status (permanent, contract, temporary). These are useful controls since technical and managerial workers are more likely to receive bonus-based compensation and shop floor workers piece-based compensation, and differences in employee composition could affect firm performance independently of the compensation structure. In addition to the quantitative variables directly provided in the dataset, I code 16 dummies for product varieties in which at least 10 enterprises specialize (for example, coke, iron ore, pig iron, raw steel, integrated iron and steel, finished steel) and use a miscellaneous category to capture the remaining varieties.

The principal defect of the CSILWS dataset is that it does not contain data on the capital stock. These data are difficult to obtain and I was only able to gather information for 228 firms, which are either larger firms or firms located in the provinces of Sichuan, Jilin, Henan, Jiangxi, or Shandong. Since reliable data on investment and capital goods prices are unavailable, I rely on book values of the capital stock net of depreciation to measure differences in capital intensity across firms. I compile data on the capital stock from other internal government sources. These sources also allow me to crosscheck the CSILWS statistics using other government sources published in different years. The sources are as follows: *Local Medium and Small Iron and Steel Enterprises Historical Financial Statistics (1983)*, *Smelting Industry Enterprise Financial Statistical Materials, Vols. 1 and 2 1985-1990 (1992)*, *Smelting Industry Enterprise Financial Statistical Materials, Vols. 1 and 2 1949-1978 (1979)*, *Reform of Small and Medium Iron and Steel Enterprises, 1979-1987 (1988)*, *Liaoning Smelting Industry Production and Labor Statistics (1962)*, *Sichuan Smelting Industry Financial and Cost Statistics 1950-1985 (1987)*, *Hunan Smelting Industry Statistical Materials,* 1949-1976 (1977), *Shaanxi Smelting Industry Statistical Materials, 1949-1977 (1978)*, *Guizhou Smelting Industry Statistical Materials, 1949-1976 (1977)*, *Jilin Smelting Industry Statistical Materials 1949-1980 (1982), Henan Smelting Industry Statistical Materials 1979-1985 (1988), Jiangxi Smelting Statistical Materials 1949-1990 (1991)*, *Shandong Smelting Industry Statistical Materials 1980-1985 (1986) and 1986-1990 (1991)*, and *Fifty Years of the Chinese Iron and Steel Industry Statistical Compendium, Vols. 1 and 2 (2003).*

With the exception of the last compendium, the data were compiled for internal government use, not for public release. The use of data from these sources has two significant advantages. Firstly, the firms under analysis constitute essentially the entire population of iron and steel firms, rather than a nonrandom sample selected by a government agency. This avoids selection bias, often a concern in previous studies which rely on samples released by the Chinese government (Li 1997). Secondly, the availability of multiple government sources dating from different time periods allows an assessment of the consistency of government data over time. In crosschecking, I find that internal government sources are generally consistent with one another, regardless of when they were compiled. Publically released plant-level data, however, such as statistics in *Fifty Years of the Chinese Iron and Steel Industry Statistical Compendium, Vols. 1 and 2 (2003)*, report some investment, capital, and employment figures for the pre-reform period which are inconsistent with internal government sources and seem implausible.[3] In addition, aggregate-level data in publically released yearbooks, while apparently accurate, do not include information on input use in 1953-1956, 1958-1961, 1963-1964, 1966-1969, and 1971-1974. Previously, no statistics have been available for the above periods, at either the plant- or aggregate-level for any industrial sector.

Collection of the data provides an opportunity to assess the relationship of productivity trends in the post-reform period to trends in earlier periods. One striking feature of the data is that radical changes in labor productivity trends occurred several times prior to the onset of enterprise reform. These breaks in trend were contemporaneous with changes in compensation policy, technology policy, and political disruptions. Since similar changes in planning policy also occurred during the late 1970s and early 1980s, this raises the question of how important these policy changes were relative to novel efforts to reform SOE management.

[3] For example, the Great Leap Forward (1958-1961) was associated with a rapid influx of agricultural labor into the iron and steel industry. Internal government sources document a ten-fold increase in the iron and steel labor force in 1958, whereas the publically released compendium does not show any jump in employment levels in this year.

A Long-Run Perspective on Policy Changes and Productivity

Studies of productivity growth in Chinese SOEs often focus on novel institutional reforms, such as rapid expansion in the use of bonuses, as a key driver of growth. A basic assumption underlying these studies is that the forces driving labor productivity growth during the reform period differed substantially from those in earlier periods. The study of reforms during the 1980s, however, should be placed in the context of changes in economic policy occurring in other periods. Many of the major policy changes which began in 1978 had precedents in policies pursued during prior decades, in particular in the early- to mid-1950s. During both of these periods, Chinese planners made use of promotions based on skills and education to motivate workers, limited the rate of employment growth in state industry, and imported foreign technologies on a large scale. These policies are likely to have had a positive effect on productivity, and could potentially be sufficient to explain why growth in labor productivity during these periods was relatively rapid.

In general, labor productivity in Chinese industry grew during periods of relative stability, such as the Postwar Recovery Period (1949-1952), the First Five Year Plan Period (1953-1957), the Post-Great Leap Forward Stabilization Period (1962-1965), and the Early Reform Period (1978-1988). However, during chaotic periods associated with extreme leftist movements, such as the Great Leap Forward (1958-1961) and the Cultural Revolution (1966-1976), productivity collapsed, erasing earlier achievements. Since extreme leftist policy makers opposed incentive pay, incentive use also fell during these periods. In Figure 1, I depict these trends in the Chinese iron and steel industry, graphing both aggregate log labor productivity and the percentage of wages distributed as bonuses and piece rates. Several major shocks to productivity are evident; sudden collapses in productivity occurred in 1958 and 1967, and recoveries occurred in 1962 and 1977. Though incentives were also applied during other stable periods, the extent of incentive use occurring during the 1980s was unprecedented. The very low level of incentive use during the Cultural Revolution reflects the aversion of extreme leftist leaders to material incentives.

Leftist movements had major impacts on industry independent of their influence on bonus use. They profoundly influenced technological borrowing, the growth rate of industrial employment, and the availability of wage increases through promotions. During the First Five Year Plan (FFYP), China purchased advanced iron and steel technologies from the Soviet Union on a large scale (Clark 1973). Throughout Chinese industry, most machinery installed during the 1950s was imported from the Soviet Bloc.[4] In 1956, Sino-Soviet relations began to deteriorate, and after 1960 the Soviets withdrew all technological assistance. To a significant extent, the collapse in technological cooperation reflected Soviet disapproval of the extreme leftist approach to development adopted by Chinese policy makers in the late 1950s. During the 1960s and early 1970s, China relied almost exclusively on domestic capital equipment. Large-scale technological borrowing resumed at the onset of the reform period in 1978, this time with the US, Japan, and Europe as the key source countries. Like shifts in the use of incentives,

[4] For aggregate statistics on the installation of capital equipment and the import of machinery by source country, see *Machinery and Electronics Historical Statistics 1949-1984* (1985).

changes in technology imports correlate closely with productivity trends. Periods of rapid borrowing during the early to mid-1950s and after 1978 coincided with accelerated productivity growth.

A second major influence of politics on industry were the shifting objectives of the central government with respect to wage and employment growth. Faced with a limited agricultural surplus, planners had to decide whether to allocate funds to increases in employment or to the promotion of skilled workers. Extreme leftist policy makers tended to emphasize rapid employment growth and wage restraint, while rightist policy makers advocated a closer adherence to the strategy recommended by Soviet advisors, emphasizing capital-intensive technologies and the provision of skill- and education-based wage increases. From 1953 to 1957, the heyday of Soviet influence, the absorption of rural labor was held to a modest pace and industrial workers were offered the opportunity to earn significant wage increases through skill-based promotions. With the onset of the Great Leap Forward in 1958, leftist policy makers rejected Soviet advice, transferring a significant fraction of the rural population into industrial occupations. In the iron and steel industry, aggregate employment jumped from 400,000 in 1957 to 4 million in 1958, returning to a sustainable level of 700,000 only in 1962. Excessive iron and steel employment had a large, negative effect on food security (Li and Yang 2005). During the next three years, the government returned to offering workers promotion opportunities and avoided increases in the industrial labor force. In 1966, the onset of the Cultural Revolution marked a return to the leftist emphasis on industrial employment growth. To decrease the likelihood of an agricultural crisis, increases in the industrial labor force were accompanied by very strict controls on industrial wages. From 1966 until 1976, promotions essentially did not occur, so that new entrants to the labor force had little incentive to acquire skills. At the onset of the reform period in 1978, promotions resumed and renewed emphasis was placed on skill and education as criteria (Shirk 1981).

Analysis of performance and political changes prior to 1978 raise significant questions for the interpretation of SOE reforms during the 1980s. Political stability, access to technology, and the use of promotions are all likely to have played key roles in the relatively successful performance of SOEs during the early-to-mid 1950s and the mid-1960s. One important question is whether positive developments in these areas can also explain improved performance during the 1980s. In particular, it is not clear whether labor productivity growth resumed during the 1980s because SOE workers and managers were provided with bonuses, or because of increases in technological borrowing, the use of promotions, and enhanced political stability. Though a statistical treatment of these issues is beyond the scope of this paper, they should be kept in mind when interpreting the results related to incentives. The key point here is that the use of the Cultural Revolution as a reference period to evaluate the effects of post-1978 reforms is problematic. Many disruptive events occurred during the Cultural Revolution, and the cessation of performance-based pay was likely a comparatively minor issue.

## The Specification and Results

In the long run, change in labor productivity depends primarily on technological progress and the accumulation of appropriate capital goods. Whether rewards for current performance accelerate productivity growth depends on whether they strengthen incentives to undertake these activities. If offering bonuses for short-run performance induces a neglect of long-term activities, then incentive use could have a negative effect on future labor productivity. To test this, I regress the log gross output per worker in 1980 prices at firm *i* at time $t$, $Y_{it}$, on the 'bonus share', defined as the fraction of wage compensation distributed on a bonus basis, $b_{it}$. To capture long-run effects, I include a 10 period series of lags of the bonus share. I also include the share of wages distributed on a piece basis, $p_{it}$, to control for other changes in incentive use.[5] The specification is listed below, where $\mu_i$ is a firm fixed effect, $\gamma_{vt}$ is a product-specific time dummy, $Z_{it}$ is a vector of additional firm characteristics (described in a footnote)[6], $P_{jt}$={$W_{jt}$, $Y_{jt}$} is a vector describing the characteristics of province j (it includes the province-level mean log wage, $W_{jt}$, and the province-level mean log labor productivity $Y_{jt}$)[7], $\nu_{it}$ captures performance determinants observed by planners but not in the data, and $\varepsilon_{it}$ is measurement error:

$$Y_{it} = \mu_i + \sum_{T=0}^{10} \beta_{1,T} b_{i(t-T)} + \beta_2 p_{it} + \gamma_{vt} + \delta_1 Z'_{it} + \delta_2 P'_{jt} + \nu_{it} + \varepsilon_{it} \quad (6)$$

Here, the coefficients of interest are the series {$\beta_{1,0}, \dots \beta_{1,T}$} on the current and lagged bonus share. In equation (6), $b_{i(t-T)}$ is influenced by whether the firm meets performance standards at time $(t-T)$, and is thus an endogenous variable. In all likelihood, bonus use is positively correlated with unobserved performance determinants, and thus the coefficient on the current period bonus share, $\beta_{1,0}$, will be biased upwards. An overestimate of $\beta_{1,0}$ will also bias other coefficients, complicating estimation of the entire vector {$\beta_{1,0}, \dots \beta_{1,10}$}. Controlling for endogeneity is thus essential to obtaining consistent estimates from this regression.

I approach the endogeneity problem using three separate methods, all of which yield qualitatively similar results. In each case, I find that bonuses had a weak positive effect on current performance and a large negative effect on future performance, and thus that the overall effect of bonus use on long-run performance was negative. I begin by using a strategy which relies on a set of assumptions to infer unobserved performance determinants from choices made by planners. The basic assumption is that planners in different provinces adopt different bonus setting practices, and thus that bonus levels in otherwise identical firms differ based on location. Cross-provincial policy differences allow the direct effects of bonuses to be separated from their correlation with unobserved productivity determinants. The second method exploits a source of predetermined variation in incentive use, early adoption of incentive pay in 1978. Until at least the mid-1980s, early adopters used consistently more bonus and piece rate payments than firms which adopted incentives later. Using a matching estimator, I compare changes in labor productivity over time among early adopters relative to firms which adopted incentives at a later date. Unlike regression analysis, matching includes effects directly due to bonuses as well as indirect effects which operate through the influence of bonuses on the evolution of other

[5] I do not focus on piece rates because they accounted for less than one-third of incentive-based pay, and were applied selectively based on their suitability to the enterprise's production process, rather than to all firms as was ultimately the case with bonuses.

[6] $Z_{it}$ includes a cubic function of the firm's age, a cubic function of the firm's employment level, the percentage of employees who are female, managerial workers, technical workers, service workers, trainees, contract workers who may be dismissed due to negligence, and temporary workers who may be dismissed at the firm's discretion. The omitted categories are production workers and permanent workers. When $Z_{it}$ is interacted with other variables, it also includes a set of dummies for 16 product varieties.

[7] Both of these averages refer to the iron and steel industry only. When calculating these averages, I exclude data from the firm under observation to avoid spurious correlation.

variables. In the third method, I return to OLS estimation, including firm-level average wages in the regression, and reincorporating observations with a bonus share equal to zero. Since planners also regulate wage increases on a performance basis, wages provide an alternative proxy for $\nu_{it}$. Using this method, I am able to control for other potential sources of productivity variation, through the inclusion of province-specific time dummies and firm-specific time trends.

In the first method, I make assumptions about bonus setting policy and the joint distribution of the shock $\nu_{it}$ and firm-level characteristics $Z_{it}$, which allow unbiased estimation of the coefficients on the current and lagged bonus shares. The assumptions allow $\nu_{it}$ to be expressed as a function of observed variables. Inserting this function into the regression specification controls for unobserved performance determinants and thus resolves the endogeneity problem. The key identification assumptions are: first, up to a constant and a slope term, planners use a consistent function to set bonuses, $f(\nu_{it}, Z_{it}, P_{jt})$; second, the joint distribution of $\nu_{it}$ and $Z_{it}$ conditional on provincial characteristics $P_{jt}$, is independent and identically distributed; and finally, the function $f(\nu_{it}, Z_{it}, P_{jt})$ is strictly increasing in $\nu_{it}$. This last assumption requires the exclusion of observations for which the bonus share is equal to zero, and accordingly I drop these from the analysis. I denote the conditional variance of the function $f\left(\nu_{it}, Z_{it}, P_{jt} | P_{jt}\right)$ as $\sigma_f^2(P_{jt})$. I assume that province- and time-specific variation enters the model through two exogenous parameters which specify provincial bonus policy: a mean bonus level, $b_{Jt}$ , and a performance sensitivity, $\alpha_{Jt}$. The bonus share is determined as follows:

$$b_{it}(\nu_{it}, Z_{it}, P_{jt}) = b_{jt} + \alpha_{Jt} f\left(\nu_{it}, Z_{it}, P_{jt}\right) \qquad (7)$$

I set the mean of $\alpha_{Jt} f\left(\nu_{it}, Z_{it}, P_{jt}\right)$ to zero by incorporating it into $b_{jt}$. Under these assumptions, $f\left(\nu_{it}, Z_{it}, P_{jt}\right)$ can be inferred from observable variables up to a constant of proportionality. To see this, note that

$$b_{it}(\nu_{it}, Z_{it}, P_{jt}) - b_{Jt} = \alpha_{Jt} f\left(\nu_{it}, Z_{it}, P_{jt}\right) \quad Var\left(\alpha_{Jt} f\left(\nu_{it}, Z_{it}, P_{jt}\right) | P_{jt}\right) = \alpha_{jt}^2 \sigma_f^2(P_{jt}) \quad (8)$$

$$\frac{f(\nu_{it}, Z_{it}, P_{jt})}{\sigma_f(P_{jt})} = \frac{b_{it}(\nu_{it}, Z_{it}, P_{jt}) - b_{Jt}}{\sqrt{Var\left(b_{it}(\nu_{it}, Z_{it}, P_{jt}) - b_{Jt} | P_{jt}\right)}} \qquad (9)$$

I define an estimate of $\frac{f(\nu_{it}, Z_{it}, P_{jt})}{\sigma_f(P_{jt})}$ as the relative bonus share, $b_{it}^r$, and estimate it as shown in equation (10), where $\overline{b_{Jt}}$ is an estimate of $b_{Jt}$, $s_{Jt}$ is an estimate of $\alpha_{jt}\sigma_f\left(P_{jt}\right)$, $\{j_t\}$ is the set of firms in province j at time t, and $J_t$ is the number of firms in $\{j_t\}$:

$$b_{it}^r\left(\nu_{it}, Z_{it}, P_{jt}\right) \equiv \frac{b_{it} - \overline{b_{Jt}}}{s_{Jt}}, \text{ where } \overline{b_{Jt}} \equiv \frac{1}{J_t} \sum_{i \in \{j_t\}, b_{it} > 0} b_{it} \text{ and } s_{Jt} = \sqrt{\frac{\sum_{it \in \{j_t\}} \left(b_{it} - \overline{b_{Jt}}\right)^2}{J_t - 1}} \quad (10)$$

Since the function $f\left(\nu_{it}, Z_{it}, P_{jt}\right) = \sigma_f(P_{jt}) b_{it}^r\left(\nu_{it}, Z_{it}, P_{jt}\right)$ is strictly increasing in $\nu_{it}$, an inverse function exists, $\nu_{it}(Z_{it}, P_{jt}, b_{it}^r) = g^{-1}(\sigma_f(P_{jt}) b_{it}^r\left(\nu_{it}, Z_{it}, P_{jt}\right))$. I approximate $g^{-1}()$ as a piece-wise linear function of the relative bonus share with five knots, $(x_0, \dots, x_4)$, and assume that firm and provincial characteristics, $Z_{it}$ and $P_{jt}$, and interactions between $b^r_{it}$ and $Z_{it}$ and $P_{jt}$ enter the function linearly.[8] The resulting approximation, $\widetilde{\nu_{it}}$, is shown in equation (11), where $1(b^r_{it} \leq x_0)$ is an indicator function equal to one if the inequality holds and zero otherwise.

$$\widetilde{\nu_{it}}(Z_{it}, P_{jt}, b^r_{it}) = \rho_0 b^r_{it} 1(b^r_{it} \leq x_0) + \ \dots + \rho_4 b^r_{it} 1(x_4 < b^r_{it}) + \tau_0' Z_{it} + \tau_1' P_{jt} + \tau_2' Z_{it} b^r_{it} + +\tau_3' P_{jt} b^r_{it} \quad (11)$$

The coefficients in this function,$(\rho_0, \dots, \rho_4)$, and the four vectors of coefficients, $(\tau_0', \tau_1', \tau_2', \tau_3')$, approximate the performance shock, $\nu_{it}$, on the basis of the observed relative bonus share, $b^r_{it}$, and a set of observed firm and province-level characteristics, $Z_{it}$ and $P_{jt}$.

Inserting the expression for $\widetilde{\nu_{it}}$ into the regression specification given in equation (6) yields:

$$Y_{it} = \mu_i + \sum_{T=0}^{10} \beta_{1,T} b_{i(t-T)} + \beta_2 p_{it} + \gamma_{vt} + \delta_1 Z'_{it} + \delta_2 P'_{jt} + \widetilde{\nu_{it}}(Z_{it}, P_{jt}, b^r_{it}) + \varepsilon_{it} \quad (12)$$

It is important to understand how the coefficients on the current period bonus share, $\beta_{1,0}$, are identified separately from $\widetilde{\nu_{it}}(Z_{it}, P_{jt}, b^r_{it})$ in equation (12). Identification requires the exclusion of provincial time dummies from the set of provincial characteristics $P_{jt}$. Under this exclusion restriction, identification comes from variation in $\alpha_{Jt}$, the sensitivity of the bonus level to the performance. Assuming that $\beta_{1,0}$ is positive, labor productivity should vary more in provinces that employ a varied range of bonus shares; that is, provinces which have a large $\alpha_{Jt}$. A similar argument applies for variation across time. Within-province variation in the bonus share increases over time, and, if $\beta_{1,0}$ is positive, the specification implies that this leads to increasing within-province variance in labor productivity.

Since the estimation listed in equation (12) contains a large number of variables, I report only a few of the coefficients. In particular, I report $\beta_{1,0}$ and $\beta_{1,1}$, the coefficients on the current bonus share and its one period lag, but I summarize results for the eight lags prior using the estimated sum of the coefficients, $\beta_{1,2} + \dots + \beta_{1,8}$, together with its associated standard error. I also report the coefficient on the share of wages distributed as piece rates, another important form of incentive pay. To provide a summary measure of the overall effects of bonuses, I report the estimated aggregate effect of current and past bonus use on log labor productivity in 1988, $\Delta\widehat{Y_b}$, together with its standard error. I estimate $\Delta\widehat{Y_b}$ using the sum of all coefficients on the bonus share, weighting the coefficients by the current (1988) and lagged (1978-1987) means of the bonus share.

Estimations of the model that include only the current bonus share suggest that it was positively correlated with labor productivity, but that the relationship was driven primarily by reverse causation. In the first specification, I compute the model including only the current bonus share, and omitting the controls for endogeneity, $\widetilde{\nu_{it}}\left(Z_{it}, P_{jt}, b^r_{it}\right)$, from the regression. This is similar to the procedure used in Groves et al. (1994). The result, shown in Table 1, specification (1), indicates that current bonus use is positively correlated with current performance. If this is interpreted as causation, the estimate of $\beta_{1,0}$ implies a substantial positive effect on labor productivity in 1988, equal to 0.22 log points. Specification (2) includes controls for endogeneity, but is otherwise identical to specification (1). The

[8] I chose knots on basis of division of $b^r_{it}$ into five equal quintiles, that is $\Pr(b^r_{it} \leq x_0) = \dots = Pr(x_4 \leq b^r_{it}) = 0.2$.

results suggest that correlation between the current bonus share and labor productivity is primarily due to reverse causation. The estimate of $\beta_{1,0}$ in this specification is reduced by four-fold, and is not significantly different from zero.

In subsequent specifications, I focus on the effects of past bonus use on labor productivity. I find that a large, negative effect of past bonuses is present regardless of whether endogeneity is controlled for. In specification (3), I include the current bonus share and a 10 year series of lags of the bonus share, but omit controls for endogeneity. As in specification (1), failing to control for endogeneity yields a large positive estimate of the effect of bonuses on performance, and this is true for the first period lag as well. The coefficients on the second through tenth period lags are all negative and their sum is large and significant, indicating that past bonus use had a large, negative effect on current labor productivity. Specification (4) is identical to specification (3), except it includes controls for endogeneity. Here, the coefficient on the current bonus share remains positive, but is small and insignificant. The sum of coefficients on the second through tenth period lags are somewhat smaller than in specification (3), but their sum remains large and significant, and all of the coefficients are negative. The attenuation of estimated negative effects of past bonus use is not surprising. Information about current performance is reflected in $\nu_{it}$ and controlling for this information could diminish the influence of historical variables which might influence the realization of $\nu_{it}$.

The major difficulty in estimating the effect of bonuses on labor productivity is that bonus use is endogenous. One approach to establishing causality is to identify an exogenous source of variation. I focus on planners' selection of firms for experimental use of incentives in 1978 as a source of exogenous variation in incentive use. In 1978, the central government instructed provincial agencies to select some enterprises for this experimental reform, while excluding others. As a result, 40 percent of firms in the iron and steel industry adopted incentive payments for the first time in at least 10 years. The intention of the policy was to gather information on the effects of incentive use, and instructions on which enterprises to include were not explicit (Liaoning Wage Policy 1980). However, provinces tended to select relatively successful enterprises for participation because they expected that the central government desired a positive outcome (Shirk 1993). This is not a significant problem for estimation because recent performance is observed in my dataset, and I control for it as a relevant variable. Moreover, unlike provincial governments, I am interested in measuring change in performance rather than performance levels.

I estimate the relative performance of these selected firms nonparametrically using the bias-corrected matching estimator due to Imbens and Abadie (2006), and provided in a Stata algorithm in Adabie et al. (2004). The estimator can be understood as specifying two potential outcomes for each firm depending on whether the firm adopts incentives in 1978, and is thus in the treated group, or does not adopt incentives, and is thus in the untreated group. The underlying assumptions of the model are: first, after controlling for observables, selection for treatment is random; and second, the support of observable variables for the treated group is identical to the support of observable variables for the untreated group. For each firm *i*, there are two potential outcomes at time t: $Y_{it}(1)$, the outcome conditional on being in the treated group, and $Y_{it}(0)$, the outcome conditional on being in the untreated group. Firms in the treated and untreated groups are matched on the basis of similarity of a set of observed characteristics, $Z_i$, where similarity is calculated on the basis of the distance norm described in Abadie et al. (2004). The parameter of interest is the average treatment effect, or the average difference between $Y_{it}(1)$ and $Y_{it}(0)$ in the population. This is computed by matching each treated firm with a similar untreated firm, and calculating the difference in their outcome variable, $\tau_{it}$. The $\tau_{it}$ calculated from each match are averaged to generate an estimate of the average treatment effect, $\hat{\tau}_t$. The estimator can be summarized as follows:

$$\hat{\tau}_t = \frac{1}{n}\sum \tau_{it}(Z_i) \quad where \;\; (\hat{\tau}_{it}(Z_i) = Y_{it}(1) - Y_{it}(0)|Z_i = Z_i) \quad (13)$$

I measure outcome as the change in log labor productivity between year t, where $t \in \{1978, \dots, 1988\}$, and log labor productivity in 1977. The outcome variable is thus a measure of the comparative productivity growth at time t of firms which adopted incentives in 1978 relative to firms which did not. In the set of variables used to define matches, $Z_i$, I include categorical variables which are constant over time, including the firm's founding year, the firm's product type, and the province in which the firm is located, as well as the 1977 level of a set of dynamic variables related to firm labor productivity, size, compensation practices, and labor force composition.[9] The estimated effects of treatment on log labor productivity, together with 95 percent confidence intervals for the effects, are depicted in Figure 2. The figure shows that treatment had a positive effect on labor productivity from 1978 to 1983, but a negative effect on labor productivity from 1984 to 1988. The negative effects after 1985 are significant and quite large, for example -0.17 log points in 1988. This estimate is not easily compared to a linear regression estimate because of the different nature of the techniques. Regressions control for the evolution of a set of observed performance determinants over time, and thus exclude the indirect effects of treatment which operate through changes in these variables. For example, controlling for an unobserved performance shock in a regression framework may attenuate the impact of historical variables if they influence the realization of the shock. The matching estimate compares the performance of treated and untreated firms, incorporating both direct effects and indirect effects which operate through changes in other variables. The large size of the matching estimate suggests that control for endogeneity may understate the negative impact of bonuses.

In the third method, I return to OLS estimation, this time including firm-level average wages in the regression, and reincorporating observations with a bonus share equal to zero. Since planners also regulate wage increases on a performance basis, wages provide a proxy for $\nu_{it}$. I also allow for dynamic effects of wage changes to test if general changes in wages have negative long-run effects, or if this is only true for bonuses. Since the evolution of wages and bonuses is regulated through similar planning mechanisms, this serves as a form of dummy treatment, controlling for the evolution of some

[9] The variables include the firm's 1977 labor productivity, product type, location (province), average wage, number of workers, percentage of service workers, percentage of managerial workers, percentage of technical workers, percentage of trainees, percentage of temporary workers, percentage of contract workers, and percentage of females. I also include the percentage of wages distributed as incentive pay in 1975. Including the use of incentive pay in 1977 would violate the assumptions of the model because firms which used incentives in 1977 are ineligible for treatment. Firms which used incentive pay in 1975 were eligible for treatment because two of the 70 firms in this category stopped using incentives in 1976 or 1977, while one of them restarted in 1978, and is thus in the treatment group.

unobserved variable correlated with both wages and bonus use. I also include information on the capital-labor ratio to see if this influences the results. As an additional robustness check, I estimate the model in a first-differenced specification which allows a firm-specific time trend to influence performance. The model is listed in equation (14), where in addition to the variables already defined, $t_i$ is a firm-specific time trend, $\xi_{jt}$ is a province-specific time dummy, $K_{it}$ is log net capital per worker, and $W_{it}$ is the log nominal wage, and *d* indicates that a first difference is being used:

$$dY_{it} = t_i + \gamma_{vt} + \xi_{jt} + \beta_0 dK_{it} + \sum_{T=0}^{10} \beta_{1,T} db_{i(t-T)} + \beta_2 dp_{it} + \sum_{T=0}^{10} \beta_{3,T} dW_{i(t-T)} + \delta dZ'_{it} + \varepsilon_{it} \quad (14)$$

I report the coefficients for the current period and one period lag of bonus and wage changes, and the estimated sum of coefficients for the eight lag series prior to this. For one specification, I also graph the coefficients. In a variety of specifications, reported in Table 2, I find that past bonus increases predicted future declines in labor productivity. In specification (1), I compute the estimation including the current and lagged bonus shares, while omitting wages. This shows how failing to include wages results in a large coefficient on the current period bonus share. In specification (2), I include wage changes from the most recent period as an independent variable. This reduces the estimated coefficient on the bonus share to 0.28, which is less than half of the coefficient estimated in specification (1). Interestingly, controlling for wage changes makes the estimated effect of lagged bonus changes substantially more negative. The sum of coefficients on the second to tenth period lags decreases from -1.80 to -2.41. It is possible that this is due to a positive correlation between past bonus increases and firms' ability to negotiate permanent wage increases with planners. If this is true and wage increases boost performance, then omitting wage changes could lead to an underestimate of the negative effect of past bonuses. I graph the estimated coefficients on the bonus share for this regression, together with a 95 percent confidence interval in Figure 3. The effects imply an estimated impact on labor productivity in 1988 of -0.17 log points, which is significant at the five percent level. This estimated loss is equivalent in magnitude to 30 percent of the total increase in mean labor productivity between 1978, when firms began offering bonuses, and 1988, the last year included in the sample.

In specification (3), I include data on lagged wage changes in the regression to test if they have dynamic effects similar to bonuses. The coefficients on the bonus share are essentially indistinguishable from those in specification (2), with the exception that the positive coefficient on the first-period lag of the bonus share is attenuated. Unlike bonuses, wage changes do not appear to have dynamic effects. As shown in Figure 4, which graphs the coefficients on wage changes, effects of lagged wage changes vary in sign and are negligible in magnitude. This indicates that incentives offered through permanent wage changes were not associated with myopic behavior. Another possibility which could confound estimation is that firms have unobserved intrinsic characteristics which determine both their performance trends and their use of bonuses, and that these drive the results. To test whether this is the case, in specification (4) I include firm-specific time trends to allow performance to be determined by unobserved characteristics. This results in a significant increase in the magnitude of the negative effect; for example, the estimated impact of bonuses on 1988 labor productivity decreases to -0.25 log points.

A final problem is the potential for unobserved changes in capital intensity to affect results. If firms were free to choose investment levels, one would expect bonuses to decrease investment activity. However, in practice, planners set investment targets for firms, and failing to meet these can result in the loss of bonus funds. Managers hoping for large bonuses may actually choose to invest more in monetary terms, but less in terms of the unobserved effort they devote to selecting, installing, and maintaining capital plant. I test how controlling for capital intensity affects results in specification (5). Data on capital intensity are restricted to 216 firms, and to accommodate the smaller sample size I use national, rather than province and product variety specific, time dummies. Though the regression suffers from some loss in precision due to the smaller sample size, the results are similar to those in other specifications. In particular, the effect of lagged bonus payments on current performance retains the same significance level and magnitude, and the implied impact on 1988 labor productivity remains large at -0.21 log points.

To elucidate the estimated effects, I use the results from specification (2) in Table 2, to estimate how the impact of expanded bonus distributions on labor productivity changed over time. In Figure 5, I depict changes in the mean bonus share over time, together with estimates of their impact on log labor productivity. Effects are measured relative to the counterfactual that the bonus share remained fixed at its 1977 level. The figure indicates that increases in bonus distributions had a modest positive impact on labor productivity during the late 1970s and early 1980s, and a large negative impact during the late 1980s. Though I lack quantitative data to verify whether the prediction is accurate, the regressions predict that the negative effect of bonuses on labor productivity trends would strengthen in the early 1990s. Enterprise case studies suggest that iron and steel SOEs did suffer deterioration in performance during the early 1990s, and Steinfeld (1998) attributes these at least in part to their pursuit of myopic management strategies during the 1980s.

The predicted effects are consistent with the view that bonuses adversely impacted social welfare, but provided benefits for enterprise insiders. From the perspective of a social planner in 1978 who discounts the future at a rate of 0.9, the average predicted effect of bonuses issued between 1978 and 1988 would be equivalent to a reduction in current labor productivity of -0.65 log points, or the loss of about fifty percent of one year's output.[10] From the perspective of an enterprise manager, on the other hand, the likelihood of separating from the firm would lead to a more positive assessment of bonuses. According to Liu and Otsuka (2004), the average tenure among iron and steel SOE managers was about 4.5 years in 1995, corresponding to about a 22 percent probability of turnover per year.[11] Relatively high turnover compared to other countries is a product of the Chinese practice of rotating managers across firms and between enterprises and supervisory agencies. Assuming that managers

---

[10] I calculate the present value of all labor productivity changes resulting from bonuses distributed between 1978-1988, as $\Delta Y_{PV} = \sum_{t=1978}^{1988} \delta^{t-1978} \Delta_{PV,t} = \frac{1}{11} \sum_{t=1978}^{1988} \delta^{t-1978} \sum_{T=0}^{10} \delta^T \bar{b}_t \hat{\beta}_{1,T}$, where $\Delta_{PV,t} = \sum_{T=0}^{10} \delta^T \bar{b}_t \hat{\beta}_{1,T}$ is the present value of labor productivity changes resulting from bonuses distributed in period $t$, $\hat{\beta}_{1,T}$ is the coefficient on the T-period lag of the bonus share, $\bar{b}_t$ is the mean bonus share in period $t$, and $\delta$ is the discount factor.

[11] Otsuka and Liu (2002) give the average tenure of state enterprises as 5.2 years, and that of township and village enterprises as 3.7 years. The sample I use includes both types of government-owned enterprise, and I use the simple average of these numbers.

have a constant probability of job separation, place value only on events occurring within their tenure, and share the same discount rate as planners, the predicted effects from their perspective are equivalent to an increase in current period labor productivity of 0.02 log points.[12] If managers served fixed terms of predictable five-year duration and thus placed no weight on productivity changes occurring after five years, their valuation would be considerably higher at 0.06 log points. This indicates that managers could have perceived the use of bonuses as personally beneficial, even if they were fully aware of their impact on future performance. Since workers were essentially tied to their jobs until retirement, a performance decline would not have been beneficial to them ex-post. However, workers in the late 1970s and early 1980s probably doubted that incentive policy would endure for decades. Incentives introduced in the mid-1950s and mid-1960s were revoked in subsequent years (refer to Figure 1). If workers expected the 1980s incentive policy to suffer a similar fate, then it would have made sense for them to behave myopically.

## Conclusion

This paper has shown that reform of the bonus system between 1978 and 1988 had a severely negative impact on SOE labor productivity in the Chinese iron and steel industry. Unlike most empirical studies of incentives, I explicitly contrast short- and long-run effects. Here, the contrast between the short- and long-run is highly revealing, as the effects differ radically in both direction and magnitude. To avoid misleading results, researchers should attempt to differentiate between short- and long-run effects in studies of incentives. The potential for these effects to move in opposite directions may be one reason why time wages are more common than performance-based payments in market economies.

The key policy implication of these findings is that the decision to apply incentives, whether in the public sector or in private firms, should be informed by an estimate of their long-run effects. Proponents of incentive payments often argue that they align the interests of principals and agents and thus promote more efficient decision-making. In the Chinese iron and steel industry, incentives had the opposite effect, encouraging managers to pursue value-destroying strategies. In practice, this effect would have been difficult to anticipate a priori. Policy makers, observing initial performance gains, probably expected these effects to continue as incentives were strengthened, and this may have encouraged nationwide adoption of incentive use. In developed market economies, owners tend to have superior means of monitoring managers, but the problem they face in distinguishing increases in current performance from changes in firm value remains important. Excessive short-run incentives in the American financial sector have been identified as a key contributor to the current banking crisis (Cheng, Hong, and Scheinkman 2009). The difficulty in distinguishing current period from long-run effects probably played a role in the diffusion of these incentive practices.

From the perspective of Chinese economic history, the key contribution of the paper is to place the shift in trend performance of state-owned industry in the mid to late 1970s in context by comparing it to earlier changes in trends. Several studies have attempted to demonstrate the effects of market-based reform on state-owned industry by comparing productivity trends prior to 1978 to those

[12]To calculate the manager's private valuation, I replace the discount factor, $\delta$, with the product of the discount factor and unity minus the manager's probability of separation, $\delta(1-0.22)$.

afterwards, giving the pre-1978 period only cursory treatment. Explicit examination of detailed pre-1978 data suggests the link between market reform and growth in state industry be treated with more skepticism. Radical breaks in performance trends occurred prior to 1978, and they seem to be more closely associated with political disorder and accompanying policy changes than with the use of incentives. Until these prior trend breaks are understood and differentiated from the trend shift in the late 1970s, researchers should be wary of vesting institutional changes unique to the 1980s with excessive significance.

Table 1: The Effects of Past Bonus Use on Performance: The Role of Moral Hazard

| | Dependent Variable (Log Gross Output Per Worker) in 1976-1988 | | | |
|---|---|---|---|---|
| Independent Variables | (1) | (2) | (3) | (4) |
| $b_{it}$<br>Current Period Bonus | 0.91***<br>(0.10) | 0.23<br>(0.21) | 0.82***<br>(0.10) | 0.09<br>(0.22) |
| $b_{i(t-1)}$<br>One Period Lag of Bonus | | | 0.29***<br>(0.08) | 0.32***<br>(0.08) |
| $\sum_{m=0}^{8} b_{i(t-2-m)}$:<br>Two to Ten Period Lag of Bonus | | | -2.04***<br>(0.49) | -1.48***<br>(0.52) |
| $p_{it}$<br>Current Period Piece-Rate | 0.28*** (0.05) | 0.29***<br>(0.05) | 0.30***<br>(0.05) | 0.32***<br>(0.05) |
| Implied Effect of Current and Lagged Bonus Shares on 1988 Labor Productivity, $\Delta\widehat{Y_b}$ | 0.22*** (0.03) | 0.05<br>(0.05) | 0.02<br>(0.07) | -0.08<br>(0.10) |
| Controls for Endogeneity | No | Yes | No | Yes |
| Fixed Effects | Yes | Yes | Yes | Yes |
| Time Dummies | Variety-Specific | Variety-Specific | Variety-Specific | Variety-Specific |
| Number of Firms | 774 | 774 | 758 | 758 |
| N | 5927 | 5927 | 5850 | 5850 |

Standard Errors, reported in parentheses, are clustered at the firm level.
***1% significance **5% significance *10% significance

Table 1 shows that the relationship between bonus use and labor productivity is primarily driven by reverse causation. The effect of current bonuses is biased upwards due to an endogeneity problem. This leads to large estimates of the effect of current bonus use in specifications 1 and 3 which do not control for endogeneity. After controlling for endogeneity, current bonuses appear to have a very small effect (as in specifications 2 and 4). The negative effect of past bonus use, on the other hand, is robust to control for endogeneity. In all of the regressions, lagged bonus payments have a large and negative effect on current labor productivity.

Table 2: The Effects of Past Bonus Use on Performance: The Role of Moral Hazard

| | Dependent Variable (Change in Log Gross Output Per Worker) in 1976-1988 | | | | |
|---|---|---|---|---|---|
| Independent Variables | (1) | (2) | (3) | (4) | (5) |
| $db_{it}$ Current Period Bonus | 0.68*** (0.08) | 0.28*** (0.08) | 0.26*** (0.08) | 0.21*** (0.08) | 0.31 (0.21) |
| $db_{i(t-1)}$ One Period Lag of Bonus | 0.23*** (0.07) | 0.18*** (0.07) | 0.13* (0.07) | 0.11 (0.08) | 0.06 (0.14) |
| $\sum_{m=0}^{8} db_{i(t-2-m)}$: Two to Ten Period Lag of Bonus | -1.80*** (0.50) | -2.41*** (0.49) | -2.42*** (0.50) | -2.53*** (0.81) | -2.47*** (0.87) |
| $dp_{it}$ Current Period Piece-Rate | 0.21*** (0.05) | 0.08* (0.06) | 0.08* (0.05) | 0.06 (0.05) | 0.11 (0.09) |
| $dW_{it}$ Current Period Wages | | 0.71*** (0.06) | 0.74*** (0.06) | 0.72*** (0.06) | 0.70*** (0.12) |
| $dW_{i(t-1)}$ One Period Lag of Wages | | | 0.10** (0.05) | | |
| $\sum_{m=0}^{8} dW_{i(t-2-k)}$ Two to Ten Period Lag of Wages | | | -0.11 (0.21) | | |
| $dK_{it}$ Capital-Labor Ratio | | | | | 0.07** (0.04) |
| Implied Effect of Bonus Share Coefficients in 1988, $\Delta\widehat{Y_b}$ | -0.03 (0.08) | -0.17** (0.07) | -0.18** (0.07) | -0.25** (0.11) | -0.21* (0.12) |
| Firm-Specific Time Trend | No | No | No | Yes | No |
| Time Dummies | Province- and Variety-Specific | Province- and Variety-Specific | Province- and Variety-Specific | Province- and Variety-Specific | National |
| Number of Firms | 768 | 768 | 768 | 768 | 213 |
| N | 7130 | 7130 | 7130 | 7130 | 1416 |

Standard Errors, reported in parentheses, are clustered at the firm level.
***1% significance **5% significance *10% significance
Table 2 shows a series of regressions which demonstrate that past use of bonuses had a negative effect on future performance. In specifications 2 to 5, I include wages to proxy for firms' unobserved productivity. This prevents correlation between current period bonus payouts and productivity from upwardly biasing the estimated effect of bonus payments. In all of the regressions, lagged bonus payments have a large and negative effect on current labor productivity. While current bonus payments have a positive effect, this effect is relatively small. As indicated by the estimated impact of current and past changes in the bonus share on 1988 labor productivity, the cumulative effect of bonuses on long-run performance is clearly negative in every specification.

Figure 1: Aggregate Data on Labor Productivity and Incentive Use in the Iron and Steel Industry

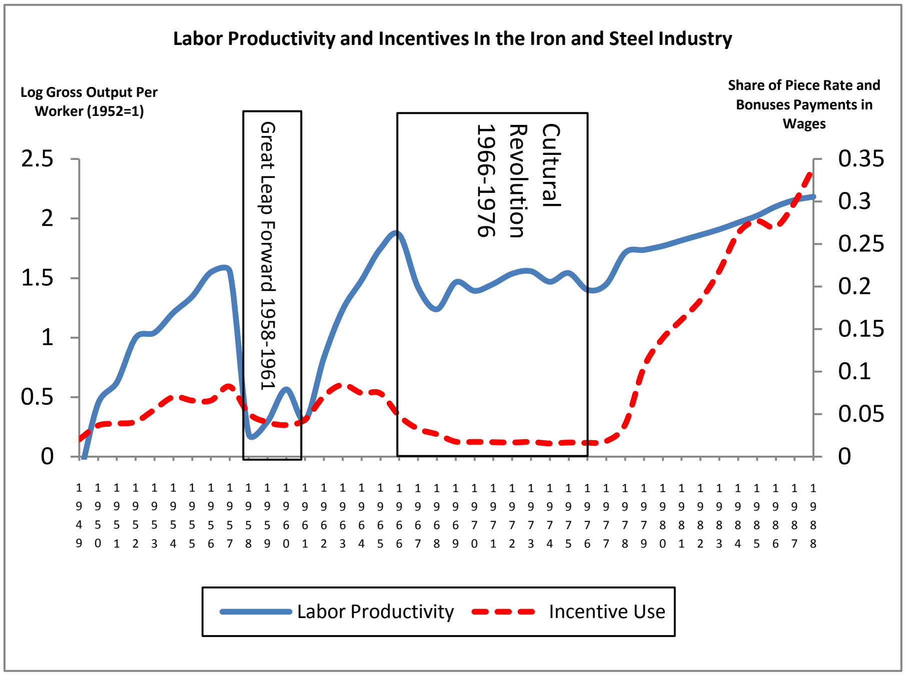


Figure 1 shows long-run changes in labor productivity and incentive use in the Chinese iron and steel industry. Periods of rising labor productivity tend to coincide with rises in incentive use. However, this may be an artifact of correlation with major political movements. The 'leftist' political faction which controlled economic policy during the Great Leap Forward and the Cultural Revolution rejected the use of material incentives. However, as I describe in the text, planners during these periods also rejected the use of education and ability as wage-setting criteria, ceased imports of foreign capital goods, and prioritized employment growth over productivity. All of these changes probably had a negative impact on labor productivity during the Great Leap Forward and the Cultural Revolution. In the post-Cultural Revolution period, it is not clear whether the rising trend in labor productivity is related to the expanded use of incentives or policy improvements in other areas.

Figure 2: Matching Estimates of the Long-run Effect of Adopting Incentives in 1978

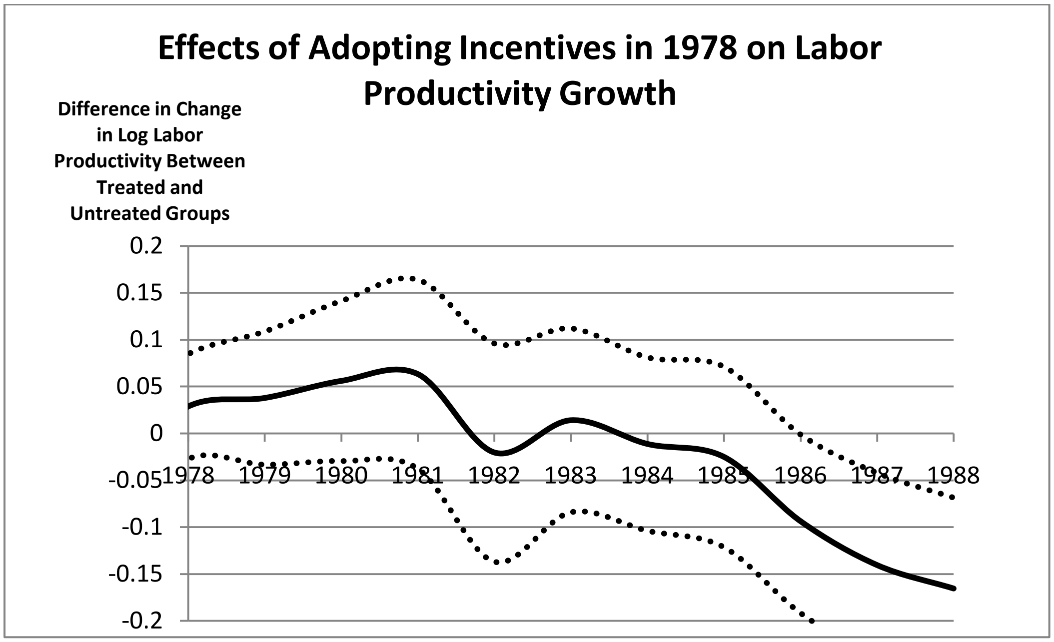


Figure 2 shows results from a matching estimator comparing labor productivity growth in firms which adopted incentives in 1978 to those which did not. The solid line shows the difference in cumulative labor productivity growth between the adopting group and the non-adopting group. When calculating labor productivity growth, I use 1977 as a base year. The dotted lines bound 95 percent confidence intervals for the estimated effect. The adoption of incentives in 1978 was associated with faster growth in labor productivity in the first four years after onset of treatment, as indicated by the positive slope in the line. After 1981, the relative performance of the treated group began to decline, reaching par around 1984, and falling -0.17 log points below par by 1988. The results thus suggest that incentive adoption had a negative effect on long-run performance.

Figure 3: Dynamic Effect of Bonus Payments on Labor Productivity

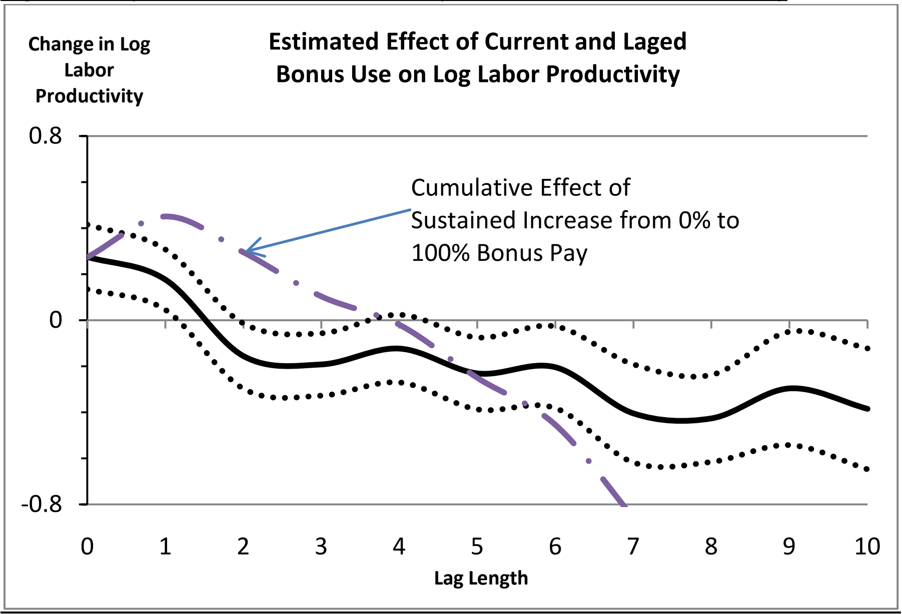


Figure 3 shows the coefficients on the current and lagged changes in the bonus share estimated in Table 2, specification 2. The solid line shows the coefficient estimates themselves and the two dotted lines depict a 95 percent confidence interval for the estimates. The dot-dash line shows the implied cumulative effect of shifting from zero percent bonus pay to 100 percent bonus pay in year 0, and sustaining this level of bonus pay in all subsequent years.

Figure 4: Absence of Dynamic Effect of Wage Changes on Labor Productivity

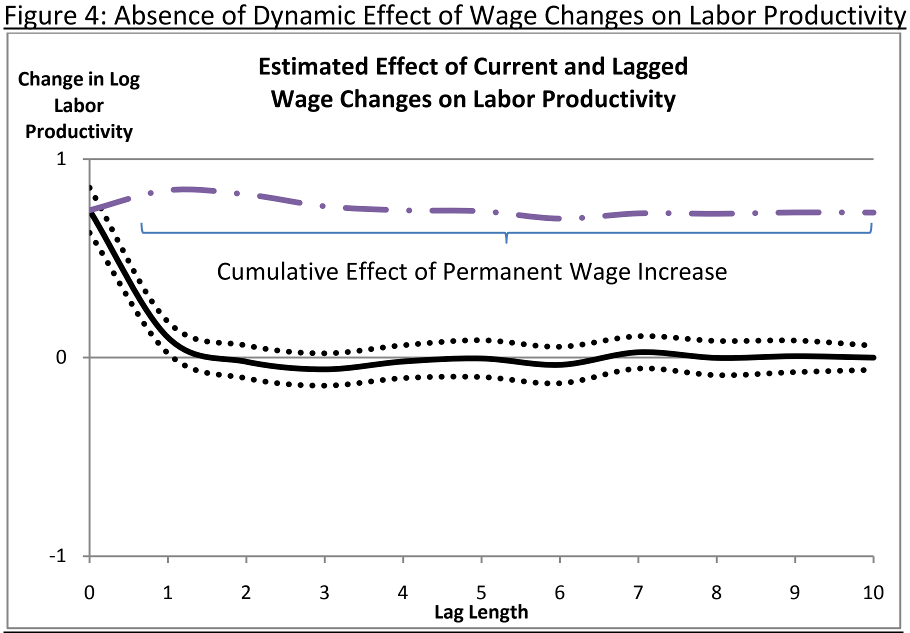


Figure 4 shows the coefficients on the current and lagged log wage changes estimated in Table 2, specification 3. The solid line shows the coefficient estimates themselves and the two dotted lines depict a 95 percent confidence interval for the estimates. The dot-dash line shows the implied cumulative effect of a permanent doubling in log wages in year 0. Unlike for bonuses, past wage increases are not associated with declines in current labor productivity.

Figure 5: Overall Effect of The Expansion of Bonus Payments on Labor Productivity

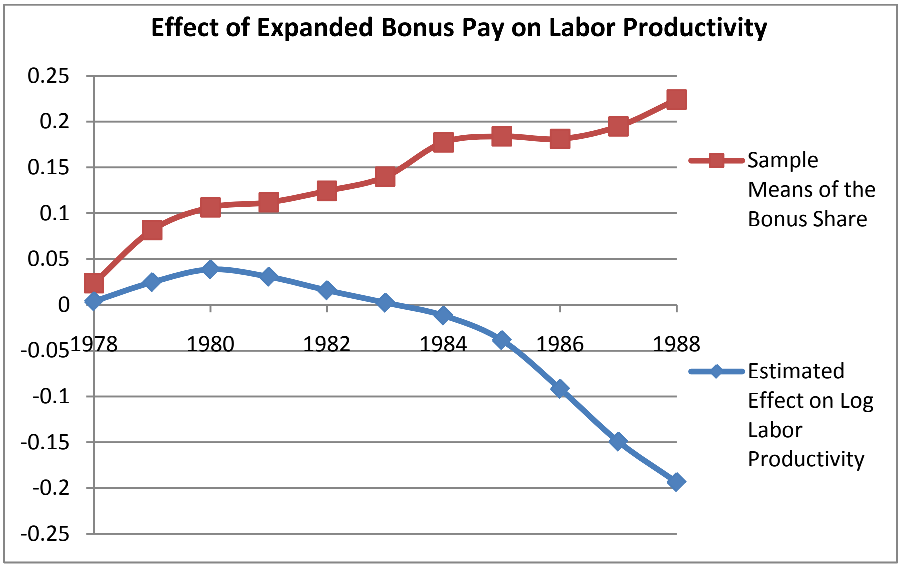


Figure 5 shows the sample mean of the bonus share in each year from 1978 to 1988 together with the estimated dynamic effect of these increases on labor productivity. I use the coefficients estimated in specification 2 in Table 2 to estimate the dynamic effect on productivity. The calculation is relative to the counterfactual that bonuses had remained fixed at their 1977 level throughout the decade. The figure indicates that bonus distributions had a modest positive impact on labor productivity during the late 1970s and early 1980s, and a large negative impact during the late 1980s. The time course of effects is similar to the matching estimates depicted in Figure 2. As in Figure 2, effects are positive prior to 1982, and negative afterwards.

Bibliography

*Data Sources*


Chinese Smelting Industry Labor and Wage Statistics, 1949-1988 Vols. 1-3. (Zhongguo Yejingongye Laodong Tongji, 1949-1988) Ministry of Metallurgy Personnel Office, 1990

Fifty Years of the Chinese Iron and Steel Industry Statistical Compendium Vols. 1 and 2 (Zhongguo Gangtie Gongye Wushinian Shuzi Huibian), Iron and Steel Industry Publishing, 2003

Guizhou Smelting Statistical Materials 1949-1976 (Guizhou Yejin Tongji Ziliao), Guizhou Province Smelting Industrial Bureau, 1977

Henan Smelting Industry Statistical Materials, 1979-1985 (Henan Sheng Yejin Gongye Caiwu Chengben Tongji Ziliao), Henan Bureau of Smelting and Construction Materials Financial Affairs Office, 1988

Hunan Smelting Industry Statistical Materials, 1949-1976 (Hunan Sheng Yejin Gongye Caiwu Chengben Tongji Ziliao), Hunan Bureau of Smelting Industry Management Bureau, 1977

Jiangxi Smelting Statistical Materials 1949-1990 (Jiangxi Yejin Tongji Ziliao), Jiangxi Province Smelting Industrial Bureau, 1991

Jilin Smelting Industry Statistical Materials (Jilin Yejin Gongye Caiwu Chengben Ziliao Huibian), 1949-1980, 1982

Local Medium and Small Iron and Steel Enterprises Historical Financial Statistics. (Difang Zhongxiao Gangtie Qiye Tongjiziliao Huibian) Ministry of Metallurgy Iron and Steel Production Technical Office, 1983

Liaoning Smelting Industry Production and Wage Statistics. (Liaoning Sheng Yejin Gongye Shengchan Laodong Tongji) Liaoning Smelting Industry Management Bureau, 1963

Machinery and Electronics Industry Historical Statistics, 1949-1984 (Quanguo Jixiedianzi Gongye Sanshiwunian Tongjiziliao Huibian), National Planning Committee for the Machinery and Electronics Industry (Guojia Jihua Weiyuanhui Jixie dianzi gongye), 1985

Reform of Small and Medium Iron and Steel Enterprises, 1979-1987 (Zhongxiao Gangtie Qieye Gaige), 1988

Shaanxi Smelting Statistical Materials 1949-1977 (Shaanxi Yejin Tongji Ziliao), Shaanxi Province Smelting Industrial Bureau, 1978

Shandong Smelting Industry Statistical Materials, 1980-1985. (Shandong sheng Yejin Gongye Tongji Ziliao Huibian) , Shandong Province Smelting Industry Corporation (Shandongsheng Yejin Gongye Zonggongsi), 1986

Shandong Smelting Industry Statistical Materials, 1986-1990. (Shandong sheng Yejin Gongye Tongji Ziliao Huibian) , Shandong Province Smelting Industry Corporation (Shandongsheng Yejin Gongye Zonggongsi), 1991

Sichuan Smelting Industry Financial and Cost Statistics, 1950-1985 (Sichuan Yejin Gongye Caiwu Chengben Ziliao Huibian), Sichuan Province Smelting Industry Bureau Financial Affairs Office (Sichuan Sheng Yejin gongyeting Caiwuchu), 1987

Smelting Industry Enterprise Financial Statistical Materials 1949-1978, Vols. 1 and 2 (Yejingongye Qiye Caiwu Chengben Tongjiziliao, 1986-1990) Ministry of Metallurgy, 1979

Smelting Industry Enterprise Financial Statistical Materials 1986-1990, Vols. 1 and 2 (Yejingongye Qiye Caiwu Chengben Tongjiziliao, 1986-1990) Ministry of Metallurgy Economic Adjustment Office (Yejingongyebu Jingji Tiaojiesi), 1992

*Published Materials*

Abadie, Alberto and Guido Imbens "Large Sample Properties of Matching Estimators for Average Treatment Effects." *Econometrica* 74, No.1 (2006): 235-267

Abadie, Alberto; David Drukker; Jane Leber Herr; Guido W. Imbens. "Implementing Matching Estimators for Average Treatment Effects in Stata." *The Stata Journal* 4, no. 3 (2004): 290-311

Asch, Beth. "Do Incentives Matter? The Case of Navy Recruiters." *Industrial Labor Relations Review* 43 (1990): 89-107

Benabou, Roland and Jean Tirole. "Intrinsic and Extrinsic Motivation" *Review of Economic Studies* 70 (2003): 489-520

Berliner, Joseph S. "A Problem in Soviet Business Administration." *Administrative Science Quarterly* 1, No. 1 (1956): 86-101

Besley, Timothy and Maitreesh Ghatak. "Competition and Incentives with Motivated Agents." *American Economic Review* 95, no. 3 (2005): 616-636

Clark, M. Gardner. *Development of China's Steel Industry and Soviet Technical Aid*. Ithaca: New York State School of Labor Relations, Cornell University, 1973

Groves, Theodore; Yongmiao Hong, John McMillan, and Barry Naughton. "Autonomy and Incentives in Chinese State Enterprises." *Quarterly Journal of Economics* 109, no. 1 (1994): 183-209

Holmstrom, Bengt and Paul Milgrom. "Multitask Principal Agent Analyses: Incentive Contracts, Asset Ownership and Job Design." *Journal of Law, Economics, and Organization* 7 (1991): 24-52

Huang, Yasheng. "Web of Interests and Patterns of Behaviour of Chinese Local Economic Bureaucracies and Enterprises during Reforms." *The China Quarterly* 123, (1990): 431-458

Kerr, Steven. "On the Folly of Rewarding A, While Hoping for B." *Academy of Management Journal* 18, (1975): 769-783

Koyima, Ryutaro. "Japanese Firms, Chinese Firms: Problems of Economic Reform in China, Part 1." *Journal of the Japanese and International Economies* 1, No. 1 (1987): 31-61

Lazear, Edward P. "Performance Pay and Productivity." *American Economic Review* 90, no. 5 (2000): 1346-1361

Li, Wei. The Impact of Economic Reform on the Performance of Chinese State Enterprises, 1980-1989. *Journal of Political Economy* 105, No. 5 (1997): 1080-1106

Li, Wei and Dennis Tao Yang. "The Great Leap Forward: Anatomy of a Central Planning Disaster." *Journal of Political Economy* 113, no. 4 (2005): 840-877

Li, Weiye and Louis Putterman. "Reforming China's SOEs: An Overview." *Comparative Economic Studies* 50, (2008): 353-380

Liaoning (No Author). *Wage Policy Compendium Vols. 1 and 2* (Gongzi Zhengce Xuanbian), Liaoning Province Labor Bureau (Liaoning Sheng Laodongju), 1980

Light Industry Economic Systems Research Group. *(Qinggongyebu Jingji Tizhi Diayanzu) "*A Report of the Situation in Sichuan Province Light Industrial Enterprises at Pilot Sites with Expanded Autonomy" *(Sichuansheng Qinggongyeqiye Kuangda zizhuquan shidian de qingkuangbaogao* ). *Economic Research (Jingji Yanjiu),* no. 2 (1980)

Liu, Kin Cheung and Xing Liu. "Management Myopia: A By-product of the Business Contract System in China?" *Asia Pacific Journal of Management* 11, no. 1 (1994): 125-131

Liu, Deqiang and Keijiro Otsuka. "A Comparison of Management Incentives, Abilities, and Efficiency between SOEs and TVEs: The Case of the Iron and Steel Industry in China." *Economic Development and Cultural Change* 52, no. 4 (2004): 759-780

O'Brian, Kevin J. "Bargaining Success of Chinese Factories." *The China Quaterly* 132, (1992): 1086-1100

Oyer, Paul. "The Effects of Sales Incentives on Business Seasonality." *Quaterly Journal of Economics* 113, no. 1 (1998): 149-186

Prendergast, Canice. "Intrinsic Motivation and Incentives." *American Economic Review* 98, no. 2 (2008): 201-205

Qian, Yingyi. "Enterprise Reform in China: Agency Problems and Political Control." *Economics of Transition* 4, no. 2 (1996): 427-447

Shearer, Bruce. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment." *Review of Economic Studies* 71, (2004): 513-534

Shirk, Susan. "Recent Chinese Labour Policies and the Transformation of Industrial Organization in China." *The China Quarterly* 88, (1981): 575-593

Shirk, Susan. *The Political Logic of Economic Reform in China.* Los Angeles: University of California Press, 1993

Steinfeld, Edward S. *Forging Reform in China: The Fate of State-Owned Industry*. New York: Cambridge University Press, 1998

Yao, Shujie. "Profit Sharing, Bonus Payment, and Productivity: A Case Study of Chinese State-Owned Enterprises." *Journal of Comparative Economics* 24, no. 3 (1997): 281-296