



The limits of meritocracy: Screening bureaucrats under imperfect verifiability[☆]



Juan Carlos Suárez Serrato^a, Xiao Yu Wang^{a,*}, Shuang Zhang^b

^a Department of Economics, Duke University, and NBER, USA

^b Department of Economics, University of Colorado Boulder, USA

ARTICLE INFO

JEL classification:

D23
D73
D86
M12
M51
O12
O15
O53
P23
P26
P48

ABSTRACT

Does bureaucratic ability predict promotion in governments? We show that self-reported performance in enforcing the One Child Policy predicts mayoral promotion in China. However, misreporting handicaps screening—a non-manipulated performance measure does not predict promotion. We show that this is consistent with a model where a government has a meritocratic objective but underestimates the imperfect verifiability of performance, rather than a model where a government is only interested in the illusion of meritocracy. Thus, despite meritocratic intentions, we challenge the notion that a successful promotion system effectively substituted for democratic institutions in explaining Chinese growth.

Bureaucracies play a central role in providing goods and services to people around the world, and the selection of bureaucrats has enormous impacts on welfare and inequality (Acemoglu, 2005; Besley and Persson, 2010; Rasul and Rogger, 2016). Scholars argue that political meritocracies—incentive systems explicitly designed to identify able leaders—may supplant democratic institutions by effectively screening for bureaucratic ability.¹ We provide a new assessment of the screening power of meritocratic incentives by recognizing the fact that bureaucrats are often evaluated on the basis of imperfectly observed and unverifiable measures of performance. How effective

are political meritocracies when performance metrics are manipulable?

We answer this question by analyzing the promotion patterns of mayors responsible for the implementation of China's One Child Policy (OCP). We find that, while provincial governments used promotion incentives to screen for mayoral ability, manipulation of reported birth statistics weakened the screening power of the meritocracy to the extent that promoted mayors were not ultimately of higher ability.

Answering this question is hampered by several challenges. First, output may be rewarded with promotion even in cases where there is no

[☆] We thank the editor and two anonymous referees for very useful suggestions. We are also grateful to Douglas Almond, Kehinde Ajayi, Attila Ambrus, Abhijit Banerjee, Pat Bayer, Charlie Becker, Prashant Bharadwaj, Allan Collard-Wexler, Julie Cullen, Esther Duflo, Pascaline Dupas, Erica Field, Raquel Fernández, Fred Finan, Rob Garlick, Bob Gibbons, Gopi Goda, Cynthia Kinnan, David Lam, Danielle Li, Hugh Macartney, Matt Masten, Michael Powell, Nancy Qian, Debraj Ray, John van Reenen, Orié Shelef, Erik Snowberg, Duncan Thomas, Chris Udry, Daniel Xu, Xiaoxue Zhao, and participants in the NBER/BREAD Fall Development Meetings, the Duke Labor/Development seminar, the Colegio de México seminar, the 2015 World Congress of the Econometric Society, the 2015 Society for the Advancement of Economic Theory, NEUDC Brown, the Stanford GSB PE seminar, the NYU Development seminar, and the MIT Sloan Organizational Economics seminar for helpful comments. This paper was previously circulated under the title “The One Child Policy and Promotion of Mayors in China.” Suárez Serrato and Zhang thank the Stanford Institute for Economic Policy Research (SIEPR) for financial support. Suárez Serrato gratefully acknowledges support from the Kauffman Foundation. We thank Pawel Charasz, Stephanie Karol, Han Gao, Matt Panhans, and Victor Ye for providing excellent research assistance.

* Corresponding author.

E-mail addresses: jc@jcsuarez.com (J.C. Suárez Serrato), xy.wang@duke.edu (X.Y. Wang), shuang.zhang@colorado.edu (S. Zhang).

¹ Bell (2015) defines political meritocracy as “the idea that a political system is designed with the aim of selecting political leaders with above average ability to make morally informed political judgments.” Bell and Li (2012) provide an overview of meritocracy in the Chinese political system from the point of view of political science.

screening motive. This underscores the need to first establish whether a set of incentives is driven by an objective to identify high-ability leaders. Second, unobserved factors that can lead to cadre advancement (e.g. political connections, unobserved ability, economic growth) may be correlated with observed performance and may confound the relation between performance and compensation. Third, it is often not possible to detect manipulation in performance metrics, which is essential to evaluating the screening efficacy of any incentive system.

We overcome these challenges by combining theoretical insights with institutional features of Chinese mayoral promotion. First, we formalize the incentive design of a government that compensates mayors who differ in privately-known ability on the basis of manipulable performance measures. The model captures our empirical setting and connects the equilibrium incentive structure with the strength of the government's screening motive (i.e. meritocracy). We show that screening is costly for the principal, in the sense that incentives that prioritize screening reduce total output. Intuitively, a principal that prioritizes screening over production will provide incentives that separate the effort exerted by agents of differing abilities, as this increases the chance that agents with higher output are also the ones with higher ability. However, these incentives reduce the effort of lower-ability agents as well as total production. We characterize this trade-off between extracting information and maximizing production, and derive testable comparative statics describing how the impact of increased output on the probability of being promoted varies with respect to features of the environment, such as promotion competitiveness and output variability. These theoretical insights enable us to establish empirically whether observed incentives are consistent with a screening (meritocratic) motive.

We apply these insights to the promotion of Chinese mayors and provide empirical evidence that provincial governments used the implementation of the OCP to screen for high-ability cadres. We focus on the OCP as it is broadly recognized as one of the top three priorities in cadre evaluation (Birney, 2014), and as this setting provides a unique opportunity to evaluate whether imperfect verifiability limited the system's effectiveness at screening bureaucrats. First, we find that mayors with better self-reported OCP performance are more likely to be promoted.² On average, reducing the reported rate of natural increase by 1 in 1000 people increases the probability of promotion by 13% and is equivalent, in its effect on promotion, to an 10% increase in GDP.³ Second, we find that promotion patterns are consistent with the comparative statics of a system of meritocratic incentives. We show that OCP performance is more predictive of promotion in provinces where OCP is a more informative signal of mayoral ability. We also show that OCP performance is less predictive of promotion in provinces where promotions are more competitive, as it is more difficult for any given mayor to achieve the maximum output—and thus be promoted—when the pool of competitors is larger.⁴ These results suggest that, beyond implementing the OCP,

² Throughout, we use the demographic concept of the *rate of natural increase*, which is defined as the crude birth rate minus the crude death rate. Targets for the rate of natural increase are set every five years in *Five-year Plans*. We measure mayors' performance in implementing the OCP by the gap between the centrally-set target of natural increase and the rate they report achieving.

³ In particular, we show that OCP matters for promotion beyond meeting the target, providing empirical evidence for Wong (2012), who reports that officials are penalized in promotion for poor population control, regardless of their performance in other categories.

⁴ These results are robust to alternative measures of noisiness and competitiveness. We use three measures of noisiness: the province-level variance in birth rates, the province-level variance in gross migration, and the average of both measures. We use the average tenure of provincial officials as a first measure of competitiveness. Intuitively, provinces with higher average tenure are more competitive, since this means slots for promotion open up less frequently for aspiring mayors. We also use the average promotion rate as a second measure of competitiveness as well as the average of both measures.

promotion incentives were used deliberately to screen for high-ability mayors.

Finally, we study whether the meritocracy succeeded at identifying high-ability mayors. Since mayors are evaluated on the basis of self-reported data, mayors may misreport. We use population auditing surveys and employ a retrospective measure of performance based on census microdata to analyze potential cheating behavior. We find that mayors adjust their manipulation of data on the rate of natural increase. This is consistent with our model's prediction that mayoral effort increases and misreporting decreases when audits increase the probability of detection.

Our model shows that the scope for manipulation limits the power of the screening mechanism to select high-ability mayors: the expected ability of the promoted mayor is a combination of the expected ability under a random promotion scheme and a scheme where OCP performance is perfectly verifiable. Our results show that provincial leaders attempted to use promotions as a meritocratic screening device, and that mayors respond to these incentives by working harder in years when audits are stronger. While provincial governments were not able to evaluate the extent of misreporting in real time, our retrospective analysis shows that, in practice, mayor manipulation rendered the meritocracy ineffective as actual OCP performance using census data is not predictive of promotion.

We explore the robustness of our results to a number of potential concerns. We assuage concerns that our conclusions are the spurious byproduct of unobserved measures of performance or political connections (Shih et al., 2012; Jia et al., 2015). By focusing on the comparative statics of the model, the finding that OCP incentives are consistent with a meritocratic objective relies on non-trivial patterns that are unlikely to result from spurious correlations. In particular, any story of potential confounders that may explain the effect of OCP performance on promotion must also account for the comparative statics on noise and competitiveness.⁵ Nonetheless, we pursue several alternative strategies to allay these concerns. First, we use the panel structure of our data and include mayor fixed effects.⁶ Second, we control for alternative, non-manipulable measures of economic growth.⁷ Third, we pursue an identification strategy to isolate variation in OCP performance that is exogenous to time-varying political connections. We use changes in the targets set by the central government in five-year plans as an instrument for OCP performance.⁸ Finally, we show the robustness of our results to other potential concerns, including measurement error, prefecture-level incentives, and alternative specifications. In all cases, we find that OCP performance has economically and statistically significant effects on promotion, that the comparative statics are consistent with a system of meritocratic incentives, and that actual OCP performance is not predictive of promotion.

Rather than attributing these results to corruption or to an aversion to meritocracy in the Chinese government, our results show that

⁵ We also derive alternative comparative statics on noise and competitiveness whenever promotion is based on an unobserved metric that is correlated with OCP performance. We show that if OCP performance is correlated with connections or with an unobservable measure of performance, we would not expect to see the empirical comparative statics that are revealed by the data.

⁶ Our regressions also include prefecture fixed effects, year fixed effects, prefecture-year characteristics, and potential determinants of the targets.

⁷ These measures include the log of economic output, economic growth rate, and non-manipulable measures of output such as railway cargo volume and electricity usage.

⁸ We show that the instrument is not correlated with economic growth, indicating that the effect of OCP performance on promotion is not a spurious byproduct of economic development. We also provide evidence for the validity of this strategy by showing that new province targets are not set to favor particular mayors, and that mayors are not strategically promoted to help them avoid toughened targets.

provincial governments attempted to use the OCP to identify and promote high-ability leaders. However, the imperfect verifiability of birth rates limited the efficacy of the meritocracy and resulted in a promotion rule that effectively promoted mayors at random. Counter to common belief (e.g., Wong, 2012), the use of the OCP as a screening tool may have lowered the human cost of the policy, since incentives that did not prioritize screening would have resulted in higher overall effort in controlling population growth. Overall, our results form a counterpoint to the argument that the success of the Chinese authoritarian government can be attributed to a successful system of meritocratic promotions (Bell and Li, 2012).

Our model makes three innovations that are at the core of our results. First, in contrast to the career concerns model of Holmstrom (1999), we analyze a setting where multiple agents with private abilities compete. This feature gives rise to the trade-off between signal separation versus production.⁹ Second, our model extends the classic literature on tournaments (e.g. Lazear and Rosen, 1981) by enriching the principal's objective to allow for an unobserved value for meritocracy. The possibility that the government values the screening motive beyond ability's role in production is crucial for studying political meritocracy. Finally, we allow for imperfectly-verifiable performance metrics to show how misreporting limits the screening power of meritocracies. These theoretical innovations open the door to studying a variety of new questions. In our setting, they enable us to use observed patterns of compensation to learn about the underlying objectives of the government, and to evaluate how effective the government is at achieving those objectives.

Our model simplifies some features of the environment that are not critical for our purposes. First, we focus on a single performance measure, in order to isolate the trade-off between production and signal extraction, while allowing for a manipulable measure of performance. In contrast, the standard multi-tasking model is single principal-single agent, and there is no misreporting of output by agents. Strategic misreporting from multiple mayors on multiple tasks complicates the analysis beyond the scope of this paper, leads us away from our main research questions, and is unlikely to affect our main results.¹⁰ Similarly, we do not directly model the role of political connections. Instead, our model derives testable predictions for the case where the measure of performance is related to an unobserved margin, such as a political connection. We find that the predictions go against the case of meritocracy, which implies that we may empirically separate the two cases. Regardless of these abstractions, our model shows that treating the structure of incentives as an equilibrium object and understanding the underlying objectives that determine them is essential to a complete understanding of meritocracy, its effect on incentives, and its efficacy in selection.

⁹ Our focus on the relation between government objectives and equilibrium incentives is more closely tied to research on organizational forms (e.g., Maskin et al., 2000) than to the career concerns literature. In contrast to the current literature that expands on the career concerns model of Holmstrom (1999) (see, e.g., Persson and Zhuravskaya, 2015, Jia, 2014, and Jia et al., 2015), this paper allows for strategic interaction between agents in their effort choice and misreporting behavior, which is crucial for analyzing the principal's objective. For example, Jia et al. (2015) analyze single-agent promotion where the agent is non-strategic—her ability is exogenously-given and she takes no actions. In this setting, there is no role for the principal to use incentives, like promotion, to influence agents' efforts, and it is not possible to infer the principal's objective from equilibrium incentives.

¹⁰ We know from Holmstrom and Milgrom (1991) that, in multi-tasking settings, noisiness in one performance measure may depress the power of the optimal, linear wage for all tasks. Even if this multi-tasking intuition generalizes to multiple agents and imperfectly verifiable output, our main result relating the principal's trade-off between signal extraction and production should remain unchanged. The key insight that the principal moves from a piece rate to a tournament as her value of screening grows should not change since the power of both schemes would be reduced.

Given the wide-ranging policies that are implemented through promotion incentives in China, the economics literature is just starting to recognize the importance of this mechanism for the effective implementation of economic policy.¹¹ In particular, the literature on political selection is not able to empirically distinguish incentives designed for screening from those designed for output production, or to identify screening made ineffective by misreporting. This paper provides a fundamental reassessment of the empirical relation between promotion and reported performance by combining a rich empirical environment with a model that is suited to the empirical setting. Our model shows that identifying a truly meritocratic objective requires tracing out the principal's trade-off between screening for ability and maximizing output, and that meritocratic policies may fail to select high ability mayors in the presence of misreporting.

While there is evidence to suggest that the implementation of the One Child Policy was used to screen mayoral ability beyond other measures such as GDP (e.g. Birney (2014); Scharping (2003); Wong (2012)), this is ultimately an empirical question. We test the rich predictions from our model and show that the implementation of the world's largest population control program in history was significantly affected by the use of promotion incentives.¹² We provide evidence both that performance measures are manipulated,¹³ and that this weakens the efficacy of the screening mechanism in identifying high-ability individuals.¹⁴

The enforcement of the OCP is uniquely suited to our analysis for several reasons. First, the central government audits population figures by conducting fertility surveys.¹⁵ The timing of the audits generates variation in audit probabilities, which are a central part of our model. Second, data from auxiliary population surveys can be used to measure misreporting and to evaluate the screening efficacy of promotion incentives. Third, changes in directives from five-year plans generate identifying variation in incentives. Finally, enforcing the OCP is recognized as the most difficult task among local officials (Scharping, 2003); thus, an official's proficiency in controlling births is likely an effective signal of a mayor's ability. In addition, we deepen the understanding of the OCP (e.g., Qian, 2009; Ebenstein, 2010; Wei and Zhang, 2011; Choukhmane et al., 2014). Our findings suggest a mechanism for the persistence of the OCP—in addition to caring inherently about reducing births, the meritocracy-minded government, unable to accurately evaluate the strength of its audits, erroneously believed that reported success at implementing the OCP was informative of mayors' ability.

¹¹ See, for example, Shih et al., (2012), Li and Zhou (2005), Zheng et al., (2014), Landry (2008), and Fisman and Wang (2015).

¹² Our results thus add to the literature on incentive design by taking theory to a unique dataset in an interesting Chinese setting. While papers in contract theory explore incentive structures given noisy signals of effort or ability (Lazear and Rosen, 1981; Baker et al., 1994; Holmstrom, 1999; Rochet and Stole, 2003), the body of empirical evidence testing these theoretical mechanisms is relatively small (Baker et al., 1994; Prendergast, 1999; Chiappori and Salanie, 2000).

¹³ Our paper is also related to forensic economics, a literature which uncovers evidence of hidden behaviors and corruption, and which studies the role of audits in limiting corruption. See Zitzewitz (2012) for a review of recent papers in forensic economics, as well as studies that use audits to detect corruption, including Olken (2006), Ferraz and Finan (2008), Ferraz and Finan (2011), and Bobonis et al. (2013).

¹⁴ Whether screening mechanisms successfully select high-ability agents for promotion is still an open question in many areas. For example, recent studies find empirical evidence of incentive distortions, including teacher manipulation of test scores (Jacob and Levitt, 2003) and of student composition in test-taking pools (Cullen and Reback, 2006). However, these studies do not investigate the impact that manipulation has on the quality of teachers hired.

¹⁵ In contrast, it is difficult to evaluate the accuracy of local GDP figures and no formal auditing is used by upper-level government to verify these reported measures of performance.

The rest of the paper is organized as follows. Section 1 describes institutional details of the OCP and its implementation. Section 2 develops a model of optimal compensation for mayors and discusses testable predictions. We describe our data in Section 3. Section 4 tests the theoretical predictions on the promotion rule and Section 5 provides evidence of data manipulation in response to audit risk and analyzes the efficacy of the promotion mechanism in selecting high-ability mayors. Section 6 concludes by discussing the role of manipulation in analyzing equilibrium promotion rules.

1. Institutional details of the One Child Policy

In 1979, soon after China's Cultural Revolution and after a decade-long economic crisis, Deng Xiaoping expressed the fear that "without birth planning, economic growth will be consumed by population growth." Since then, all economic planning has presupposed success in population control. At the national level, a specific target on population growth was set so that the total population would not exceed 1.2 billion in 2000. Chinese scientists working for the government further developed a projection that showed that, in order to achieve the population target, the optimal fertility level should be one child per woman (Scharping, 2003). This recommendation was incorporated into the family planning policy in the same year and the policy was thereafter known in the West as the One Child Policy (OCP).¹⁶

Under the OCP, a limit of one child per family was strictly enforced in urban areas, and second-child permits were issued for special exemptions in rural areas and for ethnic minorities. Some other exemptions were also granted, for example, to couples who were disabled or who lived in remote areas. Provinces with a tight policy restricted themselves to common norms for exemptions, while regions with a more relaxed policy may include other criteria. The national policy was relaxed in 1984 to allow rural couples to have a second child if their first-born was a girl.

1.1. Enforcement mechanisms

A variety of birth control methods have been used to enforce the OCP. Sterilization and insertion of an intrauterine device (IUD) after the first or second birth were implemented on a large scale. Between 1979 and 1999, the percentage of women of reproductive age who underwent sterilization rose from 21% to 35% (Scharping, 2003). Meanwhile, coerced abortions of unauthorized pregnancies have been used as a "remedial measure making up for contraceptive failures." For above-quota births, financial sanctions are the main instrument for enforcing the OCP. Depending on the location and time period, the birth of an extra child can cost a family 10%–25% of their annual income for 7–14 years. Other punishments widely used include denial of bonus payments, health and welfare benefits, denial of job promotions or even demotions in urban work units, as well as the confiscation of family farmland in rural areas.

Strong resistance and non-compliance at the grassroots level, especially in rural areas, made it very challenging to enforce the OCP. As documented by Scharping (2003), internal reports issued within the party in the 1980s and 1990s acknowledged that assaults on local birth-planning cadres were frequently provoked by coercive abortions, sterilizations, and the administration of penalties.

¹⁶ The OCP was intended to end by 2000. Amendments have relaxed this policy for single children, and it was further relaxed to a universal two-child policy in 2015. However, birth planning remains a 基本国策 (i.e., "a fundamental national policy") of the government (<http://news.sina.com.cn/c/2016-01-11/doc-ixnkky7874744.shtml>).

1.2. Mayoral promotions and the OCP

“计划生育是天底下最难的一项任务!”

“Birth planning is the hardship number one under heaven!”

A Chinese cadre (Scharping, 2003)

The central government controls the appointment, evaluation, promotion, and demotion of subnational officials in China, and the career paths of these officials are determined by the performance of their jurisdictions (Xu, 2011). The central government directly controls the key positions at the province level and grants the provincial government the power to appoint key officials at the prefecture level. The provincial government stipulates a set of performance criteria for mayors.

Economic growth, social stability, and enforcement of the One Child Policy are consistently among the highest priorities (Birney, 2014). In a published list of performance indicators of 104 prefectures in 2000, GDP per capita was used to evaluate economic growth and the birth rate was used to evaluate enforcement of the OCP (Landry, 2008). However, despite the consensus that the OCP is among the highest priorities of the Chinese government, we are not aware of empirical evidence on the role of enforcing the OCP in the evaluation of officials. This is the first study to present a bigger picture by considering the role of the OCP, in addition to economic growth, in determining promotion of officials.

A centrally-controlled planning system has monitored the local enforcement of the OCP since the 1980s. At the highest level, the State Planning Commission sets birth plan targets as part of five-year plans, with the original goal of meeting the national population goal of 1.2 billion by the year 2000. Thus, the annual province-level targets for population and birth rate are set every five years. Only national and provincial targets are set; these targets are assigned to prefectures and further distributed to lower levels (Hardee-Cleaveland and Banister, 1988). Prefectures are responsible for local implementation and submit the population and birth data to provinces. Provinces then transmit these numbers to the central government.

Birth control performance is directly linked to cadre evaluation. Hardee-Cleaveland and Banister (1988) document the following:

In Shaanxi in March 1987, the provincial party committee deputy secretary and acting governor “demanded that leaders at all levels should simultaneously grasp two kinds of production - economic production and reproduction - and take measures to do this work well and firmly. Otherwise, they are not qualified leaders.”

1.3. Data collection and (mis)reporting by Chinese mayors

Data for the evaluation of mayoral OCP performance are gathered by a birth planning commission (Scharping, 2003). The birth planning commission is in charge of workers in villages (rural), neighborhood committees (urban), and industrial enterprises (urban), who file cards on women of reproductive age, and keep track of their births. Prefectures do random checks twice a year to verify these numbers, while the birth planning commission does random checks more frequently. These local workers have no direct incentive to deviate from their assigned task, since their incomes and careers do not explicitly depend on these numbers, in contrast with the mayor. The prefecture then aggregates these local numbers and sends a prefecture-level number to province. The data manipulation (misreporting) is most likely to occur through the mayor (either directly or indirectly), since her performance evaluation depends explicitly on this number.

In addition to the birth rate, the birth planning rate, which is the percentage of total births that are authorized, is also reported from lower

level governments. In practice, lower level authorities often report very high birth planning rates of 98–99%, which are extremely unrealistic and unreliable. This is because it is much more difficult for the central government to verify whether a birth is authorized than whether a birth occurred. For this reason, the birth planning rate is not used in the evaluation of mayor performance (Scharping, 2003).

The leadership has been aware that, even for birth rate numbers, there are potential problems with data quality since the data are reported by officials whose evaluations depend on these data. Population census data are ideal as a systematic comparison with reported numbers, but are only conducted approximately every ten years (1982, 1990, and 2000).¹⁷ A mini census for 0.1% of the population was conducted in 1995. To further investigate the credibility of reported birth numbers, the State Birth-Planning Commission was charged with conducting national fertility surveys for 0.1%–0.2% of the population in 1988, 1992, and 1997. These census and national fertility surveys were organized at the province level, where they serve as the main instrument for data validation. As an example, the 1992 fertility survey uncovered an underreporting of 18% in reported birth rates. A particularly striking case of underreporting was found in Guangxi province and the leadership was forced to deliver a written self-criticism (Scharping, 2003).

Fig. 1 displays an official document from Fujian province that links OCP performance to promotion outcomes and details guidelines for local officials with respect to the implementation of the OCP. The first highlighted section states that local officials are responsible for reporting accurate birth rates and other OCP statistics. The second highlighted section states that local officials should ensure the accuracy of the reported numbers and avoid underreporting, misreporting, faking, and failing to report birth rate statistics. Finally, the third highlighted section states that the province government is responsible for investigating violations of these guidelines. If these guidelines are violated, the responsible officials are denied positive credits in their annual evaluation and their records are sent to the personnel department of the province government.

2. A tournament model with non-contractible output

We are interested in whether self-reported and actual OCP performance predict mayoral promotion, and what we can learn from any discrepancy or parity. Because we cannot observe the government's true objective, or the compensation schemes that bureaucrats actually faced, we write a model that yields a relationship between promotion patterns and the government's objectives. In particular, we show that different objectives correspond to differences in how the marginal impact of self-reported output on promotion probability varies in competition and noise. The key theoretical insight is that a tournament with one prize generates more "separation" in output across agents with different ability and convex effort costs, while an incentive structure which applies more continuous pressure, such as a piece-rate, generates more total output. The principal's relative value of screening (i.e. meritocracy) determines whether she designs incentives to magnify differences in the output of high-ability versus low-ability agents (a tournament), at the cost of total output, or to generate as much output as possible, which reduces what individual output reveals about the underlying ability of the agent (a piece-rate). We also characterize the degree to which imperfect verifiability of performance affects the efficacy of a tournament as a screening mechanism for ability.

¹⁷ Census data are collected independently by the City Bureau of Statistics (organized by the National Bureau of Statistics (NBS)). They survey every household to gather birth and population information. It is viewed as the best data on birth and population counts.

2.1. The model

The world is populated by a risk-neutral principal and $N \in \{2, 3, 4, \dots\}$ risk-neutral agents. Agents are heterogeneous only in their privately-known ability, a_i . Utility for all agents is described by $u_i(x) = x$ for $x \in \mathbf{R}$. In our setting, the principal is the provincial governor and the agents are the mayors competing for promotion.

The governor chooses a compensation scheme and which single mayor to promote in order to maximize the following weighted sum of total output produced and the expected ability of the promoted mayor:

$$\omega E \left[\sum_{i=1}^N y_i \right] + (1 - \omega) E[a_i \mid i \text{ is promoted}]. \tag{1}$$

In our setting, output y_i is the birth rate achieved by mayor a_i .¹⁸ The weight ω may capture the trade-off between output maximization and an inherent value of meritocracy, between maximizing short-run versus long-run output, and between maximizing output in one task versus maximizing output across several tasks, to name a few interpretations. If $\omega = 1$, the governor focuses only on maximizing output. If $\omega = 0$, the governor focuses only on promoting the mayor with the highest ability. ω is determined by the central government in our empirical setting and is held constant across regions.

Mayor a_i produces output $y_i = e_i + \varepsilon_i$, where e_i is unobservable and non-contractible effort, and $\varepsilon_i \sim \exp(\lambda)$ is noise with mean $E(\varepsilon) = \frac{1}{\lambda}$, variance $V(\varepsilon) = \frac{1}{\lambda^2}$, and is iid across agents.¹⁹ The effort cost to a_i is $c(e_i) = \frac{1}{a_i} \exp(e_i)$. These assumptions ensure that a given level of effort is less costly and that marginal cost of effort is lower for higher-ability mayors, that higher levels of effort are more costly for all mayors, and that effort cost increases at an increasing rate. Assume for notational convenience that $a_1 > a_2 > \dots > a_N \geq 1$; a mayor with a higher i index has lower ability. The governor and the mayors know the average ability in the population, $\bar{a} = \frac{1}{N} \sum_{i=1}^N a_i$.

Output is not observable or contractible, but the governor can compensate mayors based on their self-reports \hat{y}_i .²⁰ Suppose the governor audits mayors' reports, but these audits are limited in their scope. The governor commits to firing any mayor who can be proven to be over-reporting (the mayor incurs disutility $F < 0$ from being fired), but proving this is difficult. In particular, the audit is sufficient proof with certainty if and only if $\frac{\hat{y}_i}{y_i} > \delta$: the mayor exaggerates by a lot. If the mayor only slightly exaggerates, that is, $\frac{\hat{y}_i}{y_i} \in (1, \delta]$, then the audit is sufficient proof with only probability $p \in (0, 1)$. That is, given that a mayor is over-reporting, the audit generates sufficient proof with probability p and the mayor is fired, but generates insufficient proof with probability $(1 - p)$ and the mayor is not fired. In other words, even if a governor knows a mayor is lying, if she cannot prove it, she cannot fire him. The strength of the audit is captured both by p and δ . Both $\delta = 1$ and $p = 1$ capture the case where output is contractible.²¹

¹⁸ The principal prefers a lower birth rate, and agents must exert more effort to achieve a lower birth rate. "Higher y is better" is standard, but in the case where output is the birth rate b , we can define output $y \equiv \frac{1}{b}$.

¹⁹ Assume that $\lambda \in (0, 1)$, a parametric assumption for objects to be well-defined: this ensures that $V(\varepsilon) \in (1, \infty)$.

²⁰ See Appendix A.5 for the case where output is contractible.

²¹ Punishing liars is both an actual feature of our empirical setting (mayors are told explicitly they will incur demerits in their performance review if they are caught lying), as well as strategic for the principal. Suppose instead that the principal did not commit to punishing detected and proven liars, believing that she can correctly back out true output from reported output, because she knows the degree of inflation. Then any mayor could profitably deviate by inflating even further. This leads to unraveling, since mayors will all pool on reporting infinity.

福建省人口和计划生育工作
责任考核奖惩办法(试行)

第一条 为了稳定低生育水平,提高人口素质,确保人口和计划生育工作目标的实现,为建设海峡西岸经济区和我省全面建设小康社会创造良好的人口环境,根据《中华人民共和国人口与计划生育法》、《中共中央国务院关于加强人口与计划生育工作稳定低生育水平的决定》和《福建省人口与计划生育条例》及其他有关规定,制定本办法。

第二条 市、县(区)、乡镇(街道)党委、政府(办事处),省、市、县(区)党委、政府工作部门,村(居)党组织和村(居)民委员会、企事业单位、各社会团体和相关责任人,在人口和计划生育工作中应给予奖励或追究责任的,适用本办法。

第三条 对在人口和计划生育工作中作出显著成绩的单位和个人,给了精神和物质奖励。成绩突出的,分别对党政主要领导、分管副职和人口计生委(局);主任(局长)给予奖励,并发给奖金。

第四条 有下列情形之一的,对有关单位和相关负责人应予以追究责任:

(一)未完成年度人口和计划生育工作主要责任目标的。

市、县(区)党委、政府及计生行政部门的主要责任目标为六项指标:1、出生人口政策符合率;2、出生人口性别比;3、统计准确率;4、群众满意率;5、计划生育基本知识知晓率;6、工作经费投入。

其他单位的主要责任目标为上述指标的第1、5、6项。

(二)违反国家人口计生委规定的“七个不准”,发生重大恶性案件的;

(三)在一年内对实行计划生育的公民兑现政策规定的奖励低于当年度应奖励总量的70%的;

(四)虚报、瞒报、漏报、伪造、篡改或者拒报人口和计划生育统计数据,报表统计准确率低于95%的;

(五)在经费投入上弄虚作假,截留、挪用、克扣计划生育经费的;

(六)对违反计划生育政策的党员和干部及村(居)民委员会成员未按有关规定处理,造成恶劣影响的;

(七)有关部门没有履行计划生育综合治理职责的,或制定出台的政策规定不利于计划生育工作开展的;

(八)因工作不力,造成计划生育工作严重滑坡的;被计划生育领导小组黄牌警告或单列管理的;

(九)机关、企事业单位人员违反生育政策被发现后未及时处理;

(十)村(居)班子软弱涣散或处于瘫痪状态,又未及时调整整顿,造成计划生育工作无人管理或失控的;

(十一)人口和计划生育部门及其工作人员不依法履行管理和服务职责的。

对上述情形按如下办法认定:

出现第(一)项情形的,以上级人口和计划生育部门对其年度人口和计划生育工作目标责任制考核的结果为准。

出现第(二)至第(十一)项情形的,由上级人口和计划生育行政部门负责核实。涉及到需要进行人事调整或党纪;政纪处分的,会同组织人事部门或纪检监察机关核实。

第五条 被列为追究责任的单位,取消其当年度综合性荣誉称号和综合性评先评优资格。其应负责任的党政主要领导或法定代表人、分管领导及直接责任人,当年度考核中不能评为称职及以上等次,情节严重的要给予免职、责令辞职,并给予了相应的党纪政纪处分。对负责任人而已调离岗位的,一年内要跟踪处理。

第六条 出现第四条(一)、(六)、(八)、(十)项规定情形的,对单位和党政主要领导或法定代表人;出现第四条其它项规定情形的,对有关部门及其分管领导和直接责任人,均按第五条规定给予相应的责任追究。

第七条 人口和计划生育行政部门负责对年度人口和计划生育目标管理责任制执行情况进行考核。考核中要注重检查单位和领导干部履行职责的情况。考核结果要向同级党委、政府或人口和计划生育领导小组和上级计生行政部门报告,提出实行计划生育工作奖励或责任追究的建议。向时,向同级组织人事部门和纪检监察机关通报有关情况。

Note: This document from Fujian Province outlines guidelines for local officials on the One Child Policy and links performance to promotion outcomes. The first highlighted section states that local officials are responsible for reporting accurate birth rates and other OCP statistics. The second highlighted section states that local officials should ensure the accuracy of the reported numbers and avoid underreporting, misreporting, faking, and excluding birth rate statistics. Finally, the third highlighted section states that the province government is responsible for investigating violations of these guidelines. If they are violated, the responsible officials are denied positive credits in their annual evaluation, and their records are sent to the personnel department of the province government. Source: <http://yz.zfxxgk.gov.cn/ShowArticle.asp?ArticleID=75204>


Fig. 1. Official document from Fujian Province.

The governor chooses between two compensation schemes: a *tournament* based on reported outputs where one mayor is promoted, or a *piece rate*. Note that mayors can still be observed to be promoted while facing a piece rate on OCP performance, as their promotion can depend on their performance in non-OCP tasks.

The timing of the game is as follows:

1. Each mayor observes her own private ability a_i and chooses effort level e_i , which is non-observable and non-contractible by the governor.

2. Each mayor's output y_i is realized. Output remains private information.
3. Each mayor submits a private report of output to the principal: \hat{y}_i .
4. The governor audits each mayor. If the mayor reported $\hat{y}_i = y_i$, she is truthful and will not be wrongfully charged with lying. If the mayor reports $\hat{y}_i > y_i$, audits detect and are sufficient proof with probability 1 if $\frac{\hat{y}_i}{y_i} > \delta$, and with probability p if $\frac{\hat{y}_i}{y_i} \in (1, \delta]$. Mayors proven to be lying are fired.



Note: This figure plots equilibrium effort levels under the optimal piece rate and tournament compensation mechanisms given a set of parameter values. The figure shows that tournaments lower overall effort but increase effort separation across mayors of different abilities.

Fig. 2. Numerical example of effort separation.

5. The governor promotes one mayor based on any criterion of her choice, and compensates mayors who are not fired based on their self-reported output.²²

There are three “types” of pure strategy equilibria²³

1. “All lie”: all the mayors misreport their output.
2. “All truth”: all the mayors truthfully report their output.
3. “Partial lie”: some mayors misreport and some mayors report truthfully.

We characterize the partial lie/partial truth equilibrium in Appendix A.6. In particular, when audits are weaker, such that most or all agents misreport in equilibrium, the lowest-ability (and lowest-performing) agents have the smallest incentive to misreport. This is because the lowest-performing agents need to misreport more in order to win the tournament. In order to avoid punishment, the lowest-performing agents simply tell the truth, which decreases the likelihood that they are promoted even further. The remaining mayors are likely to be promoted and all misreport. Since the mayors that are likely to be promoted all misreport, the rest of our analysis focuses on the “all lie” equilibrium.²⁴

2.2. Equilibrium promotion patterns

We show that if the governor sufficiently values meritocracy, then tournaments are preferred over piece rates. We characterize the relationship between self-reported OCP performance and probability of promotion when the tournament is chosen and when the piece rate is chosen, as well as how this relationship varies with different degrees of

competition between mayors and different degrees of noisiness in production. Our final result characterizes the extent to which a tournament mechanism based on self-reported performance selects mayors who are actually of high ability as performance becomes less and less verifiable.

2.2.1. The relationship between compensation scheme and the government’s objective

We first characterize the compensation scheme set by the governor, depending on the relative importance of meritocracy.

Proposition 1. *There exists an $\tilde{\omega} \in (0, 1)$ such that, if $\omega > \tilde{\omega}$ (the governor values output for production relatively more than for screening), the governor uses a piece rate mechanism. If $\omega < \tilde{\omega}$ (the governor values output for screening relatively more than for production), the tournament mechanism is chosen.*

Please see Appendix A.2 for a detailed proof.²⁵

The intuition behind this result is the following. When the governor’s sole concern is output maximization, continuous incentive pressure is the cheapest way to induce total effort. The convexity of effort costs makes it less costly to smooth effort across agents of all ability levels. On the other hand, if the government’s objective is to use observed output to identify the highest-ability agent, then tournaments are preferred. Since the reward for output is discontinuous, and the marginal cost of effort is decreasing in ability, the differences between the effort exerted by lower- and higher-ability mayors will be greater than under the continuous piece rate.²⁶ This “effort separation” across the heterogeneously-able individuals reduces total output produced but enables the governor to infer mayors’ abilities from output observations. When the governor sufficiently values promoting the highest-ability mayor, the tournament mechanism becomes optimal as it induces bigger differences in effort exerted by mayors of different ability, thereby increasing the expected ability of the promoted mayor. However, the manipulability of output limits informativeness, as the

²² Note that the principal is not allowed to promote zero mayors (although she is allowed to set a bonus of zero). This would be an extra incentive tool for the principal to induce effort. Ruling this out is realistic (when a slot opens up, some agent does get promoted to fill it), and also “works against us” in the sense that this makes the tournament a less potent tool. Since one of our key goals in constructing this model is to show that there exists an intuitive subset of the parameter space, specifically, a screening motive, for which the principal’s equilibrium choice of compensation scheme is a tournament, showing that the principal chooses to use a weakened tournament when she cares about screening implies that a more flexible tournament would be even better.

²³ We focus on pure strategies as it is unclear how to empirically interpret mixed strategies in OCP enforcement effort and reporting in a static game.

²⁴ We test this assumption empirically in Table A.21 and Table A.22, which show that low-performing mayors (those who do not reach the target) misreport less and are also less likely to misreport at all.

²⁵ If we allow for a combination of tournament and piece rate, we see that a mechanism with both elements (e.g. a tournament with bonus B and piece rate payments for individual outputs) weakly dominates pure tournament or pure piece rate. However, as $\omega \rightarrow 0$, the optimal mechanism converges toward pure tournament, and as $\omega \rightarrow 1$, the optimal mechanism converges toward pure piece rate, so we can still compare tournament-heavy with piece rate-heavy mechanisms and our results continue to apply.

²⁶ Note that the governor cannot condition compensation directly on ability, as it is private information.

governor can only contract on reported output and detects lies imperfectly.²⁷

Fig. 2 clearly illustrates the intuition. Mayors are ordered by ascending ability along the x-axis. The chosen parameters are provided at the bottom of the figure. Given those parameters, we solve for the best piece rate and the best tournament bonus for the governor. The dark-grey bars indicate effort exerted under the piece rate scheme, and the light-grey bars indicate effort exerted under the tournament. Both the “effort separation” and the lower total effort generated by the tournament are evident.

2.2.2. Equilibrium properties of promotion under a pure production objective ($\omega = 1$)

When $\omega = 1$, so the governor cares about the OCP for population control and not as a means to screen mayoral ability, the governor compensates self-reported output (lower birth rate) with a piece rate. Thus, conditional on agent ability, the promotion rule is independent of OCP performance.

Proposition 2. *If $\omega = 1$, mayors are compensated with a piece rate and, conditional on ability, increasing self-reported OCP output does not increase mayor a_i 's probability of promotion. This is true regardless of the noise (λ) and the competitiveness (N) of the environment.*

Proposition 2 guides the econometrician to compare promotion outcomes for mayors with higher and lower reported output. However, the results of Proposition 2 no longer hold when this comparison is not conditional on a given agent's ability, which requires additional tests.

Corollary 1. *If $\omega = 1$ and the governor bases promotions on a task other than OCP performance that is positively correlated with ability, mayors are compensated via a piece rate on self-reported OCP performance. In addition, unconditional on ability and conditional on performance in the promotion task:*

- (i) *Self-reported OCP performance increases mayor a_i 's probability of promotion.*
- (ii) *Self-reported OCP performance has a larger effect (increase) on the probability of promotion in noisier environments (smaller λ).*
- (iii) *Self-reported OCP performance does not have a differential effect (increase) on the probability of promotion in more competitive environments (larger N).*

The proofs of Proposition 2 and Corollary 1 are presented in Appendix A.3.

Given that the governor compensates mayors with a piece rate when $\omega = 1$, it can be shown that, in noisier environments, higher-ability mayors exert differentially higher effort compared to lower-ability mayors. This is because, in noisier environments, the governor sets a lower piece rate (the slope of the wage in reported output is flatter). Lower-ability mayors decrease their effort differentially more than higher-ability mayors in response to this lower piece rate, since their marginal cost of effort is higher. Since ability is positively correlated with the dimension on which promotion is based, it must be that the impact of having higher ability on the probability of promotion in noisier environments is larger. But we know that ability is also positively correlated with output. Hence, it must be that the observed impact of increasing output on probability of promotion in noisier environments is also larger.

Conditional on performance in the task on which promotion is based, the impact of self-reported OCP performance on probability of

²⁷ This result survives allowing the governor to choose a scheme which includes both piece rate and tournament components. A scheme with both components always weakly dominates pure piece rate and pure tournament; however, as $\omega \rightarrow 1$, the pure piece rate again becomes optimal, and as $\omega \rightarrow 0$, the pure tournament again becomes optimal.

promotion cannot vary by competitiveness, since it is rewarded with a piece rate. Thus, Corollary 1 guides the econometrician who may be concerned that she is unable to separate variation in mayor's self-reported OCP performance from variation in ability.

2.2.3. Equilibrium properties of promotion under screening and production objectives ($\omega < 1$)

Proposition 3. *If $\omega < \tilde{\omega}$, mayors are compensated with a bonus in a tournament where the mayor with the highest reported OCP performance who is not caught lying is promoted. Further, in the equilibrium where all mayors misreport:*

- (i) *Self-reported OCP performance increases mayor a_i 's probability of promotion*
- (ii) *Self-reported OCP performance has a larger effect (increase) on the probability of promotion in less noisy environments (larger λ .)*
- (iii) *Self-reported OCP performance has a larger effect (increase) on the probability of promotion in less competitive environments (smaller N)*

The technical details of the proof of Proposition 3 are presented in Appendix A.4.

Result (i) is clear, since the principal has no reason to discourage higher output, beyond firing mayors who are caught lying. Result (ii) holds because output is more informative of mayor ability in less noisy environments. Thus, the governor has more confidence that the high production mayor she is promoting is indeed of higher ability. Result (iii) holds because increasing output increases the probability of achieving the maximum output by a larger amount when there are fewer competitors. That is, there are fewer outputs a mayor has to beat in order to have the maximum output in a less competitive environment.

We take the contrasting predictions of Corollary 1 and Proposition 3 to the data as they allow us to empirically differentiate between the case where promotions are driven by a meritocratic objective, and the case where, despite the lack of a meritocratic motive, the correlation between OCP performance and promotion is driven by an omitted variable.

2.2.4. The predictive power of self-reported versus actual output for promotion when output is non-contractible

Our final result characterizes the degree to which mayor manipulation decreases the screening ability of the tournament mechanism when governments wish to use OCP to learn about mayor ability.

Proposition 4. *If $\omega < \tilde{\omega}$, so that the tournament compensation is optimal in the “all lie” equilibrium with all mayors misreporting output:*

- (i) *The expected ability of the promoted mayor when output is not contractible is a weighted average of the expected ability of the promoted mayor when output is contractible and the expected ability of a randomly-drawn mayor, where the weight depends on audit strength p :*

$$E^{\text{non-contract}}[a_i | i \text{ promoted}] = \tilde{p}E^{\text{contract}}[a_i | i \text{ promoted}] + (1 - \tilde{p})\bar{a}$$

- (ii) *The weaker the audit (lower p), the closer the expected ability of the promoted mayor is to a random draw: $\tilde{p}'(p) < 0$.*
- (iii) *When audits are completely uninformative, the expected ability of the promoted mayor is exactly the population average: $\tilde{p}(0) = 0$.*

See Appendix A.4 for details.

Typically, researchers studying audits have analyzed experiments (see, e.g., Olken, 2006). In practice, the treatments raise the audit risk from near non-existent to near certainty. Proposition 4 gives us insight into interior audit risks, which are likely to be more realistic and feasible as long-run policies. This contribution of our model can be applied to questions beyond promotions in China.

3. Measuring promotion and OCP performance

The system that monitors and sets birth targets was built in the 1980s with the ultimate goal of containing population growth by year 2000. Our study focuses on the time period 1985–2000 since it is ideal for studying the relationship between mayoral promotions and the implementation of the OCP, and since systematic prefecture-level data became available starting in 1985.

3.1. Sample of mayors

We collected Chinese mayoral data by digitizing a complete list of mayors in office from 1985 to 2000 from two series of hard copy records: *City Gazetteers*, published by the gazetteer office of each city, and the *City Development Yearbook*, published by the Chinese Urban Development Research Council. The list includes the mayor's name, as well as the year and month at the start and end of her term. We obtained data on 967 mayors in 258 prefectures and 28 provinces between 1985 and 2000. While data on Chinese political leaders at the provincial level are more commonly accessible, to our knowledge there are no such comprehensive data for Chinese mayors before 2000. Landry (2008) is the only other example we know of that uses mayor data in years 1990–2000.

3.2. Promotion

Promotion is defined as an upward move in the political career. The most natural upward move for a mayor is becoming the party secretary in the same or a different prefecture, which is the definition of mayor promotion in the existing literature (Landry, 2008).²⁸ This measure, though convenient, ignores other possible moves above the prefecture level, including provincial governor or vice-governor, minister of central ministries, etc. We define a mayor as being promoted if there is an increase in her bureaucratic rank to any of the following positions at the end of her term:

1. Prefecture: party secretary in the same or a different prefecture.
2. Province: provincial governor or vice-governor, party secretary or vice-secretary, party committee member, chairman or vice-chairman of the People's Political Consultative, chairman or vice-chairman of the People's Congress.²⁹
3. Central: minister or vice-minister of central ministries.

A mayor is not promoted if she continues as mayor, moves to positions of the same bureaucratic rank, or exits politics. First, one could continue as mayor in the same or a different prefecture. In our data, forty mayors served in two prefectures. If one is transferred from the first city to the second, she is not promoted in the first city and her promotion status in the second city depends on her move after serving the second time. Second, one could be promoted to positions in the provincial government that have the same bureaucratic rank as mayor: director or vice-director of provincial departments, assistant to the provincial governor, etc. Finally, while one could leave politics by working

²⁸ In Landry (2008), promotion is defined as being promoted to the party secretary in the same prefecture or elsewhere. This definition underestimates the likelihood of promotion because mayors could move to higher-ranked positions at the province or central level. We use the most complete definition of promotion based on the bureaucratic rank. In Table 1, we show that 15% of promotion of mayors in 1985–2000 was above the prefecture level.

²⁹ An alternative is to define promotion to province-level positions based on administrative division, i.e., province is a higher administrative division than prefecture. However, this is more controversial than the definition based on bureaucratic rank. In our definition, if a mayor becomes the director of a department in the provincial government that has the same bureaucratic rank as mayor, she is not defined as being promoted. In our data, only 40 out of the 967 mayors serve in multiple prefectures.

in industry, we only observe this for three mayors in our data. See Appendix B for details on the measurement of promotion in the data.

3.3. OCP performance

Targets for the rate of natural increase are set every five years in *Five-year Plans*. The rate is the crude birth rate minus the crude death rate. We digitized the targets for the rate of natural increase from *Provincial Five-year Plans* in 1985, 1990, and 1995. For example, a province's 1985 plan sets the target for annual rate of natural increase in 1986–1990.

We use two measures of the birth rate to compare the reported OCP performance with actual performance: (1) the official rate of natural increase from published data that are reported to provincial governments, and recorded in *City Statistical Yearbooks*, and (2) a retrospective birth rate from microdata of 1990 and 2000 population censuses that are not observed by the provincial government on a yearly basis, and the rate of natural increase is measured by the crude birth rate from census minus the crude death rate in *City Statistical Yearbooks*. A mayor's reported OCP performance is measured by comparing the reported rate of natural increase with the target from the corresponding five-year plan. The lower the reported rate relative to the target, the better the mayor's OCP performance.³⁰ Unfortunately, not all prefectures publish data on the rates of natural increase consistently in 1985–2000. On average, 80% of prefectures report yearly data, except in 1988, when no prefectures published birth rate data.³¹ In our sample of mayors, 697 out of 967 are matched with the reported rates.³²

We measure actual OCP performance by the gap between the target and the rate using census data. Census data are collected independently by the City Bureau of Statistics (organized by the National Bureau of Statistics (NBS)). They survey every household to gather birth and population information.³³

³⁰ To the best of our knowledge, prefectures face a common province-level target. Our main specification controls for potential determinants of targets at the prefecture-level, including the percentage of childbearing-age women, the percentage of Han population, and the percentage of rural population. In addition, in Section 4.2.3, we estimate the effects of these determinants on province-level targets and use prefecture-level variation to compute prefecture-level targets. Our results are robust in these specifications.

³¹ Table A.1 summarizes the number of prefectures that report birth rate data by year in our analysis sample.

³² Note that when prefectures do not report yearly data, these data include birth rate, economic growth, and population statistics. The difference between our initial sample of 967 mayors and our final sample of 697 is due to changes in administrative boundaries that prevent us from linking prefectures across Censuses. Importantly, this change in sample size is not due to selected reporting by mayors.

³³ The NBS organizes a quality control survey after each census to check for unreported people (for example, hidden children). The survey sample for the 1990 census is around 170,000 people. The NBS finds 1 unreported birth per 1000 births. It is viewed as the best data on birth and population counts. An interesting set of papers evaluates the quality of the 1990 census data. There is no consensus about underreporting in the census: Banister (1992) and Johanson and Nygren (1991) argue that there is no underreporting, while Zeng et al. (1993) argues the opposite. The latter paper backs out "actual births" by counting children in the mid-1990s and accounting for deaths. They find that births of female children are underreported in the 1990 census. We do not think this is a problem for our results, since we find that mayors underreport even relative to a potentially underreported census, our estimates are a weak lower bound on cheating. We also find that the extent of cheating decreases in audit years. This result would be overturned by an underreported census only if the census is particularly underreported in audit years, which seems very unlikely. Finally, we find that promoted mayors do no better in the census than non-promoted mayors. This would be overturned by an unreported census only if the census over-reported births especially for promoted mayors.

Table 1
Summary statistics.

	(1) Mayor			(2) Mayor-year		
	Mean	Std. Dev.	Obs	Mean	Std. Dev.	Obs
<i>Promotion</i>						
Promoted	0.53	0.5	697	0.15	0.35	2058
Prefecture party secretary	0.45	0.5	697	0.12	0.33	2058
Province government or party leaders	0.06	0.24	697	0.02	0.13	2058
Central ministries	0.02	0.13	697	0	0.07	2058
Tenure (year)	3.79	2.03	697	2.76	1.86	2249
<i>Rate of natural increase (RNI, per 1000 population)</i>						
Recorded rate	7.66	3.29	697	7.63	3.71	2058
Rate from census	8.28	5.52	675	8.43	5.74	1895
Target rate	10.57	2.32	697	10.53	2.51	2058
<i>OCP performance</i>						
ReportedRNIisbelowtarget	0.8	0.31	697	0.8	0.4	2058
Target RNI - reported RNI	2.91	3.19	697	2.9	3.7	2058
RNIfromcensusisbelowtarget	0.74	0.4	675	0.74	0.44	1895
Target RNI - RNI from census	2.27	5.16	675	2.14	5.42	1895
Real GDP (million RMB)	8517	11359	697	9348	12255	2058
Log (GDP)	3.89	1.04	697	3.99	1.06	2058
<i>Prefecture-year controls</i>						
Population (1,000)	5383	49415	697	6132	114802	2058
Percentage of urban population	0.31	0.17	697	0.32	0.17	2058
Investment (million RMB)	3973	9187	697	4382	10463	2058

Notes: Please refer to Section 3 for details on data sources.

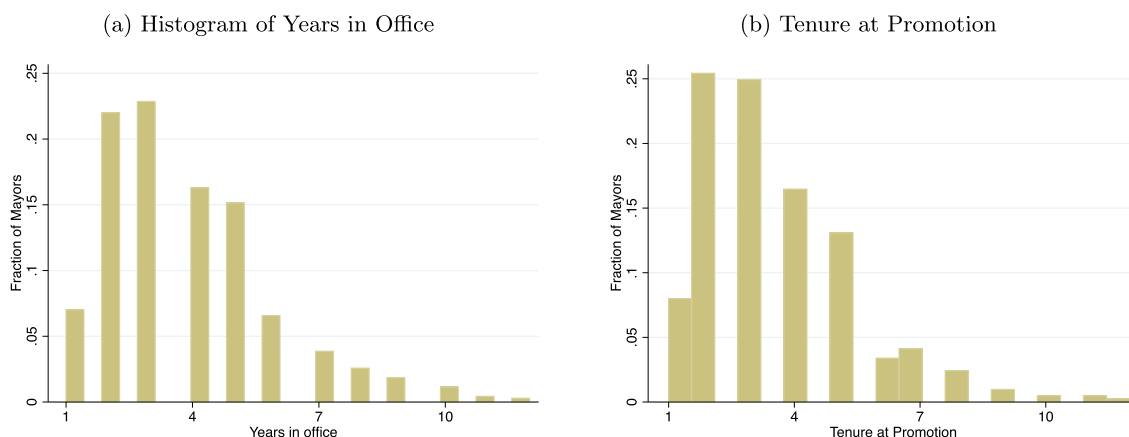


Fig. 3. Histogram of years in office.

Notes: Fig. 3a shows the distribution of years in office per mayor. Fig. 3b shows the distribution of tenure at promotion of mayors who were promoted.

We compute the crude birth rate retrospectively using microdata from the 1990 and 2000 censuses; these data are observed by provincial government only in census years. Crude birth rates in 1986–1989 come from the 1990 census and those in 1990–2000 come from the 2000 census. The main concern with census data is the potential for internal migration, since prefecture of birth is not observed for migrants in all years. Migration was tightly restricted under the *Hukou* system until its relaxation in the 1990s. Figure A.1 shows the percentage of migrants in 1982–2000 from census and population surveys. The migration rate of the entire population remained under 2% in the 1980s and slowly increased to 4% in 1995. In our estimations, we allow for year-specific effects of these average migration measures for years 1990–1994 and 1985–1990. The most significant increase in migration occurred between 1995 and 2000, with the migration ratio reaching 11% in 2000. Since we observe the prefecture of origin for migrants from 1995 to 2000, we use information on out-migration and in-migration by prefecture and year when constructing birth rate measures. We find that ignoring migration leads to very small underesti-

mation of birth rate by 0.2 (per 1000 population) in 1995–2000. In Appendix B, we discuss in greater detail how we measure the crude birth rate from census data and how we account for migration.

3.4. Summary statistics

Our main analysis sample includes 697 mayors in 211 prefectures and 28 provinces.³⁴ Table 1 reports the summary statistics at both the mayor level and mayor-year level. 53% of mayors were promoted to a higher-ranked position, with 45% promoted to party secretary in the same or a different prefecture, 6% to leadership at the province level, and 2% to central ministries. Among all promotions, 15% moved above the prefecture level, suggesting a substantial underestimation of mayor promotion by the definition in previous studies. On average, mayors spent 3.8 years in office. Fig. 3a shows the distribution of years in office.

³⁴ Column 1 of Table A.1 summarizes the number of mayors by year in our analysis sample.

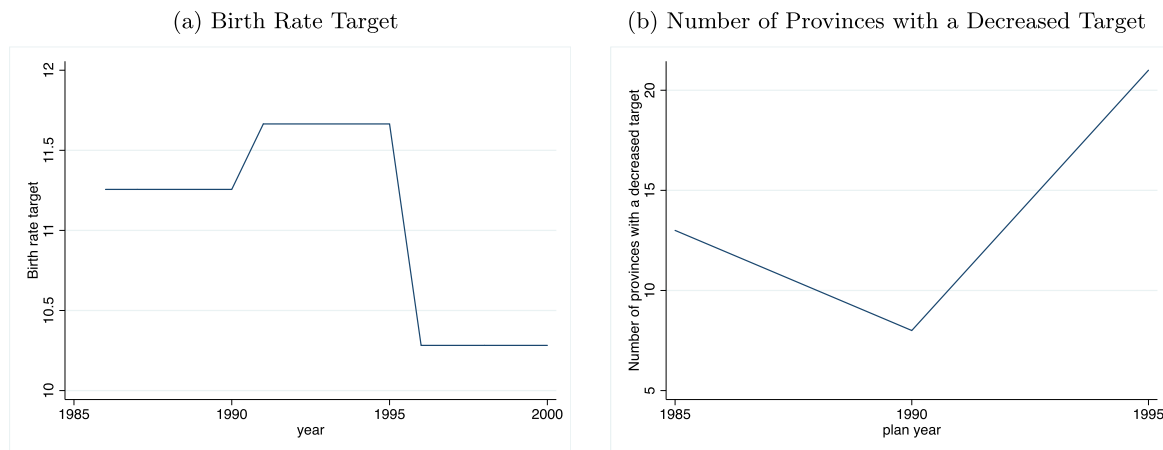


Fig. 4. Birth rate targets. Notes: Fig. 4a shows the average target for the rate of natural increase across provinces set in five-year plans. Fig. 4b presents the number of provinces that experienced a decrease in the target in each of the five-year plans in 1985, 1990 and 1995.

Most mayors were in office from two to five years. The turnover rates are especially high in the second and third years. Tenure at promotion has similar properties and is graphed in Fig. 3b.³⁵

A key variable of interest is OCP performance, which is measured by the gap between the birth rate target and a given measure of birth rate. The average reported rate of natural increase is 7.7 per 1000 people, while the rate computed using census data is 8.3. Both are lower than the target average of 10.6. On average, the reported rates are 3 births per 1000 people below target. 80% of mayors reported the rates lower than their assigned target. In comparison, the rates from census data suggest that only 74% of mayors were below their specified target. In our analysis, we use changes in targets as instruments for OCP performance. Fig. 4 presents (a) the average target across provinces, as well as (b) the number of provinces that experienced a decrease in the target in each of the five-year plans. In addition, Table A.13 reports correlations between the targets and province characteristics. This table shows targets are not set to favor mayors in more concentrated provinces or those that represent a larger share of a province’s population. Moreover, Table A.13 shows that the rate of mayoral promotion does not affect how targets are set.

Data on annual nominal GDP at the prefecture level in 1985–2000 come from *City Statistical Yearbooks*. We use the nominal GDP and national current price index (CPI) to compute real GDP. The average real GDP in the mayor-year sample is 9348 million RMB. Landry (2008) shows a published list of performance indicators of 104 prefectures in 2000, where GDP per capita is listed as the measure to evaluate economic performance of mayors. Therefore, in our main specification, we follow the official evaluation metrics and use log GDP to measure economic performance while controlling for population. Finally, we compiled prefecture-year controls from *City Statistical Yearbooks*, including population, percentage of urban population, and government investment.

4. Mayoral promotion and OCP performance

We now analyze whether reported OCP performance affects the promotion of mayors. Our first specification tests this prediction by exam-

³⁵ Our results are robust to excluding mayors that were only in service for less than three years, as we show in Table A.24. Note that the number of prefectures in a province varies from 1 to 21. Figure A.2 plots the histogram of the number of cities per province. We also show in Table A.15 that our results are robust to excluding mayors that represent large cities within a province as well as provinces where population is concentrated in few cities.

ining whether reported OCP performance increases a mayor’s probability of promotion using a linear probability model:

$$Promoted_{icpt} = \beta_1 OCP_{cpt}^{reported} + X_{icpt}\gamma + \mu_i + \eta_{cp} + \lambda_t + \varepsilon_{icpt}, \quad (2)$$

where i denotes the mayor, c the prefecture, p the province, and t the year. μ_i are mayor fixed effects, λ_t are year fixed effects, and we also include prefecture fixed effects, η_{cp} , as some mayors served two different prefectures. The dependent variable, $Promoted_{icpt}$, is a binary outcome that is equal to 1 if mayor i in prefecture c of province p is promoted in year t and 0 otherwise. X_{icpt} is a vector of time-varying attributes of mayor i or prefecture c in year t , including the mayor’s tenure, and prefecture-year log of real GDP, log of population, percentage of urban population, log of investment, and migration controls. We also control for OCP_{cpt}^{Census} in this regression. Finally, we allow for errors to be correlated at the province level.

The key regressor of interest is reported OCP performance. Superior performance in implementing the OCP corresponds to a lower reported rate of natural increase compared to the target: $OCP_{cpt}^{reported} = Target_{pt} - BirthRate_{cpt}^{reported}$. From conversations with local officials, we know that these targets have been used as a benchmark for expectation. That is, while meeting the expectation does not guarantee rewards, local officials are rewarded if they exceed the expectation. For this reason we allow $OCP_{cpt}^{reported}$ to have differential effects on promotion depending on whether the target was met (i.e., $OCP_{cpt}^{reported} > 0$).

The identifying assumption of Equation (2) is that OCP performance is uncorrelated to other factors that may drive mayor promotion. We address this concern in three ways. First, the panel nature of our data allows us to include mayor fixed effects, μ_i , which control for time-invariant ability that may affect the initial placement or political connections to province-level officials. Second, in Section 4.1 we use predictions from our model that characterize the comparative statics of OCP performance across provinces with different degrees of noisiness and competitiveness. As a third strategy, in Section 4.2 we use an instrumental variables approach that elicits variation in OCP performance using changes in targets from five-year plans. This strategy exploits variation in reported OCP performance that is uncorrelated with either changes in connections or other changes in unobserved margins.³⁶ In addition, we perform a battery of robustness checks in Section 4.2.

³⁶ In particular, this strategy assuages concerns that time-varying political connections (as in Jia, 2014) confound the effect of OCP performance on promotion.




Fig. 5. OCP performance and promotion.

Notes: In Fig. 5, the x-axis represents the residualized OCP performance and the y-axis represents the residualized promotion probability where we control for person, city, and year fixed-effects. OCP performance reported is the target for the rate of natural increase minus the reported rate of natural increase. Panel (a) shows the relation in the sub-sample where the reported rate is above target (negative OCP performance). Panel (b) plots the relation when the reported rate is equal to or below target (non-negative OCP performance). Column (6) of Table 2 reports results from a regression used to plot this figure.

We begin by showing the correlation of reported OCP performance and promotion in Fig. 5. The x-axis represents the residualized OCP performance and the y-axis represents residualized promotion probability, where we control for person, city, and year fixed-effects. Panel (a) of Fig. 5 shows the relation in the sub-sample where the reported rate of natural increase is above target (negative OCP performance); we do not observe any correlation between reported OCP performance and promotion. Panel (b) of Fig. 5 plots the relation when the reported rate is equal to or below target (non-negative OCP performance). In this case, OCP performance is positively correlated with probability of promotion. Note that while beating the OCP target increases a mayor’s probability of promotion regardless of GDP performance, these findings also show that OCP performance matters for promotion beyond simply meeting the target.³⁷

Table 2 reports estimates of Equation (2). We start by showing in column (1) that reported OCP has no effect on promotion when we look at the full sample of mayors. As we show in Fig. 5, the effect of reported OCP on promotion depends on whether or not a mayor is below the target. In columns (2)–(5), we focus our regression analysis in the sub-sample with non-negative OCP performance.³⁸ Consistent with the graphical presentation, estimates from column (2) through (5) all show that mayors with better reported OCP performance are more likely to be promoted. In this and other tables, we focus on the results from the richest specification in column (5). Decreasing the rate of natural increase by 1 per 1000 increases the chance of promotion by 1.7 percentage points, or around 13% of the probability of promotion. To gauge this magnitude, consider that the interquartile range of OCP performance, conditional on the regression model, is 1.2. Thus,

³⁷ Evaluation costs often cause principals to set targets to weed out non-contenders, such as admissions committees which ignore students with low GRE scores or hiring committees which ignore candidates with low GPAs.

³⁸ This restriction is consistent with the model in Section 2.1. In particular, recall that the lowest-performing mayors are not likely to misreport. Moreover, they are not likely to be relevant for promotion decisions. In our data, there is always at least one mayor that reaches the target for every province in every five-year-plan period.

if a mayor’s OCP performance increases from the 25th- to the 75th-percentile of the distribution, her probability of promotion increases by 16%. To compare with GDP, we also show the estimate of log GDP and find that increasing GDP by 1% increases the chance of promotion by 17.2 basis points. These estimates suggest an economically large effect of OCP performance compared to economic growth, since decreasing the rate of natural increase by 1 per 1000 is equivalent in its effect on promotion to a 10% increase in GDP. Finally, we also fit a spline in column (6) where OCP performance reported is interacted with indicators of being below the target and being above the target, respectively, controlling for the full set of controls. We find similar results that non-negative OCP performance is positively correlated with promotion.

Relative to our model, these results are consistent with Chinese governors having strong tastes for meritocracy (i.e., $\omega \leq \tilde{\omega} < 1$) since, in this case, Proposition 3 predicts that OCP performance should be predictive of promotion ($\beta_1 > 0$). However, one concern is that governors promote mayors on an unobserved measure of performance that is correlated with OCP performance. As discussed in Corollary 1, even if mayors have no taste for meritocracy ($\omega = 1$), we might find that $\beta_1 > 0$. We distinguish between these two cases by testing additional predictions across provinces with different informational and competitive environments.

4.1. Testing model implications for the promotion rule

According to Proposition 3, if governors have strong tastes for meritocracy, captured by small ω (i.e., $0 \leq \omega < \tilde{\omega} \leq 1$), we should expect to find smaller effects of $OCP_{cpt}^{reported}$ on promotion when performance is noisier, and when more mayors are competing for promotion. In contrast, Corollary 1 shows that if $\omega = 1$ and OCP performance is correlated with the unobserved metric used in promotion, we should expect to find larger effects of $OCP_{cpt}^{reported}$ on promotion when performance is noisier, but that competition should not change the effect of $OCP_{cpt}^{reported}$ on promotion. To test these predictions, we augment the linear probability model in Equation (2):

$$Promoted_{icpt} = \beta_1 OCP_{cpt}^{reported} + \beta_2 OCP_{cpt}^{reported} * Noise_p + X_{icpt}\gamma + \mu_i + \eta_{cp} + \lambda_t + \epsilon_{icpt} \tag{3}$$

$$Promoted_{icpt} = \beta_1 OCP_{cpt}^{reported} + \beta_3 OCP_{cpt}^{reported} * Comp_p + X_{icpt}\gamma + \mu_i + \eta_{cp} + \lambda_t + \epsilon_{icpt} \tag{4}$$

From these equations, a meritocratic promotion rule would imply $\beta_1 > 0$, $\beta_2 < 0$, and $\beta_3 < 0$. In contrast, if $\omega = 1$ and OCP performance is correlated with the unobserved metric used in promotion, we would expect to find $\beta_1 > 0$, $\beta_2 > 0$, and $\beta_3 = 0$.

In order to implement the test in Equation (3), we use two measures of the noisiness of OCP performance as a signal of effort and ability. The first measure is the standard deviation of gross migration (in-migration and out-migration), and the second one is the standard deviation of the rates of natural increase by province in the census data. Intuitively, a province with more gross migration or a province with a more variable rate will make it harder for a province-level official to disentangle the noise from the true performance. As these measures have no cardinal interpretation, we normalize them to have mean zero and standard deviation of one. We also use the average of these measures as a third measure.

Table 3 presents estimates of Equation (3). Column (1) reports our preferred estimate from Table 2 for comparison. Column (2) reports a negative coefficient for the interaction of reported OCP performance and the standard deviation of migration, indicating that the marginal effect of increased OCP performance on promotion is decreasing in this measure of noisiness. In columns (3) and (4), we replace the migration measure with the standard deviation of the census birth rate and

Table 2
OLS regression of promotion on reported and actual OCP performance.

	Promotion = 1					
	(1)	(2)	(3)	(4)	(5)	(6)
OCP Performance Reported	0.002 (0.004)					
OCP Performance from Census	0.000 (0.002)					
OCP Performance Reported X Below Target		0.011** (0.006)	0.016* (0.009)	0.016* (0.010)	0.017* (0.010)	0.017* (0.009)
OCP Performance from Census X Below Target		0.000 (0.003)	-0.002 (0.005)	-0.002 (0.005)	-0.002 (0.005)	-0.002 (0.004)
OCP Performance Reported X Above Target						0.005 (0.013)
OCP Performance from Census X Above Target						-0.007 (0.009)
Log GDP	-0.011 (0.009)	-0.010 (0.011)	0.146*** (0.056)	0.146*** (0.057)	0.172*** (0.054)	-0.134 (0.396)
Observations	1832	1483	1483	1483	1483	1544
R ²	0.06	0.07	0.58	0.58	0.58	0.49
Test $\beta^{Reported} > \beta^{Census}$ (p-val)	0.37	0.08	0.07	0.08	0.07	0.06
Full Sample	Y					Y
Sample of Mayors Below Target		Y	Y	Y	Y	
Year FE	Y	Y	Y	Y	Y	Y
Person FE			Y	Y	Y	Y
City FE				Y	Y	Y
Prefecture-year controls					Y	Y

Notes: OCP performance reported is the target for the rate of natural increase minus the reported rate of natural increase. OCP performance from census is birth rate target minus birth rate from census data. Tenure fixed effects and migration controls are included in all regressions. Prefecture-year controls include log(population), percentage of urban population and log(government investment). Standard errors are clustered at the province level and are bootstrapped. Table A.18 explores additional specifications of this regression. * significant at 10% level; ** significant at 5% level; *** significant at 1% level.

with the average of the two measures, respectively, and find strikingly similar results. Table 3 also presents estimates of marginal effects at different points in the distribution of our noise measures as well as the p-value of a one-sided test of the hypothesis that $\beta_2 > 0$. We find consistent results in all specifications with marginal effects that are decreasing and statistically significant at the 25th- and 50th-percentiles.

Fig. 6a plots estimates of marginal effects normalized by the average probability of promotion using estimates from column (4) for different quantiles of the distribution of average noise, with larger quantiles indicating a noisier signal. The y-axis is the predicted percentage change in the probability of promotion of increasing OCP performance by 1 per 1000. In the visual presentation, the effect of OCP performance on promotion continuously decreases as the signal becomes noisier. In provinces where the signal is the noisiest, the effect of OCP performance on promotion is null. By contrast, in provinces in the 25th-percentile of the distribution of noisiness, increasing OCP performance by 1 per 1000 leads to a 23% increase in the probability of promotion.

We use two measures of competitiveness at the province level in order to implement the test in Equation (4). The first is the proportion of mayors that are never promoted in each province during our fifteen years of data, which is equivalent to one minus the promotion rate. The second measure is the average tenure of positions above the bureaucratic rank of mayors. In provinces with longer tenure at these upper-level positions (and thus less turnover), mayors have fewer opportunities of being promoted and must work harder at proving themselves worthy of promotion.³⁹ As in the case of our measures of noisiness, we

normalize the measures of competitiveness and also consider the average of the two measures.

Table 4 presents estimates of Equation (4) using a similar layout to Table 3. Columns (2)–(4) report negative coefficients for the interactions of OCP performance with our different measures of competitiveness. The one-sided tests reject the hypotheses that $\beta_3 > 0$. Similarly, we find decreasing marginal effects that are statistically significant at lower quantiles of the competitiveness measures. Fig. 6b uses the results in column (4) to plot estimates of the percentage increase in the probability of promotion from lowering the rate of natural increase by 1 per 1000 at different quantiles of average competitiveness and finds a quantitatively similar pattern to that of Fig. 6a. Overall, we find that, in provinces where promotions are more competitive, the marginal effect of reported OCP performance on promotion is smaller, that is, $\beta_3 < 0$.

One potential concern is that our measures of noise and competitiveness are strongly correlated such that Tables 3 and 4 are not providing independent evidence. Table A.2 explores this possibility and shows that our noisiness measures are not statistically related to the competitiveness measures. Indeed, at most 4% of the variation in competitiveness can be explained by the noise measures, and vice versa. Table A.2 also shows that our measures of noise and competitiveness are not clustered in any one region and reports correlations with other demographic characteristics. Specifically, while the SD of Birth Rate is negatively correlated with urbanization and education, the SD of Migration and the competitiveness measures are not correlated with province-level demographic characteristics. In addition, Table A.3 estimates β_2 and β_3 jointly and finds similar results to Tables 3 and 4 Fig. 6c reports marginal effects as a function of both noise and competitiveness. To ease interpretation, we only report effects that are statistically significant at the 10% level.

To summarize, we test the predictions of our model for how OCP performance should affect promotion across provinces with different levels of competition and noise in the output variable. We find that $\beta_1 > 0$, $\beta_2 < 0$, and $\beta_3 < 0$, which is consistent with a meritocratic

³⁹ In order to implement this strategy, we digitized hard-copy records on the term information of all province-level officials ranked higher than mayors. The average tenure of provincial officials ranges from 3 to 6 years across provinces. The distribution of average tenure across provinces is presented in Figure A.3.

Table 3
Heterogeneous effects of reported OCP performance on promotion by signal noise.

	(1)	(2)	(3)	(4)
OCP Performance	0.017* (0.010)	0.019** (0.008)	0.018** (0.009)	0.019** (0.008)
OCP Performance X Noise (SD Migration)		−0.015* (0.008)		
OCP Performance X Noise (SD Birth Rate)			−0.011 (0.010)	
OCP Performance X Noise (Both)				−0.022* (0.011)
Log GDP	0.172*** (0.056)	0.180*** (0.054)	0.164*** (0.053)	0.170*** (0.054)
Observations	1483	1483	1483	1483
R ²	0.584	0.588	0.586	0.588
Mean of Dependent Variable	0.132	0.132	0.132	0.132
Year FE	Y	Y	Y	Y
Person FE	Y	Y	Y	Y
City FE	Y	Y	Y	Y
Prefecture-year controls	Y	Y	Y	Y
Marginal Effect at 25th pctl		.03*** (.01)	.026** (.012)	.031*** (.01)
Marginal Effect at 50th pctl		.021** (.008)	.015* (.009)	.019** (.008)
Marginal Effect at 75th pctl		.001 (.014)	.012 (.01)	.005 (.011)
One-Sided Test of Interaction (p-val)		0.031	0.124	0.025

Notes: OCP performance reported is the target for the rate of natural increase minus the reported rate of natural increase. Noise (SD Migration) is SD of gross migration (number of in-migrants + number of out-migrants) at the province-level. Noise (SD Birth Rate) is SD of retrospective birth rate in the census data at the province-level. Both noisiness measures are standardized. Noise (Both) is the average of SD Migration and SD Birth Rate. Tenure fixed effects and migration controls are included in all regressions. Prefecture-year controls include log(population), percentage of urban population and log(government investment). Standard errors are clustered at the province level and are bootstrapped. * significant at 10% level; ** significant at 5% level; *** significant at 1% level.

promotion rule, and which rules out the case where a positive effect of OCP performance on promotion is due to an omitted variable. These results show that provincial governors are instructed by the central government to value OCP performance both because population control is inherently valued and because doing so may select high-ability mayors for promotions.⁴⁰

4.2. Robustness

We explore the sensitivity of our results to a host of potential problems, including time-varying political connections as a source for endogeneity, targets that may vary at the prefecture-level, measurement error in OCP performance, as well as alternative specifications of our estimating equation.

4.2.1. Alternative measures of economic performance

Our measure of economic performance follows the official evaluation metrics for mayors in Landry (2008). We also check the robustness of our comparative statics by controlling for alternative measures of economic performance. We first control for GDP growth rate, which is commonly used in the literature on the promotion of provincial leaders, and then control for less-manipulable measures such as electricity usage and railway cargo volume. In Table A.6, we find that the heterogeneous

⁴⁰ These results also discipline alternative models. For instance, while we assume that ω is fixed across provinces, one could alternatively consider province-varying tastes for meritocracy. However, in order for this story to match our empirical results, it would have to be the case that very competitive provinces and provinces where output is a noisier measure of effort and ability are precisely the provinces that have low tastes for meritocracy, which seems unlikely.

effects of reported OCP performance on promotion by signal noise are robust to including any of these alternative economic measures. Similarly, Table A.7 shows that the heterogeneous effects of reported OCP performance on promotion by competitiveness are also robust to including alternative economic measures.

4.2.2. An instrumental variable for reported OCP performance

We now turn to an instrumental variables approach, which alleviates concerns that time-varying unobservable factors, such as expanding political networks, are biasing our results. Our strategy leverages the fact that the targets are set at the province level by the central government in five-year plans. Changes in these targets generate “surprise changes” in reported OCP performance among mayors in office. For example, if there is a decrease in the target, it is harder for mayors to get closer to the target and thus achieve a better OCP performance. We use decreases in the target to instrument for reported OCP performance.⁴¹ Fig. 4a shows the average target from 1985 to 2000. The targets were changed twice during these fifteen years; the 1990 plan saw an average increase from 1986 to 90 to 1991-95 and the 1995 plan saw an average decrease from 1991 to 95 to 1996-2000. For a given plan, however, there was substantial variation in whether a province experienced an increase or a decrease in target. Fig. 4b shows that

⁴¹ We use negative changes in targets to avoid analyzing cases where mayors are promoted by “getting lucky” through a relaxation of standards. In unreported results, we find a similar pattern when we use all changes in targets as the instrument. In order to increase the precision of the estimate, we also control for interactions between our instrument and our measures of noise and competitiveness. These controls leverage the model’s insights that the strength of the incentive depends on the competitiveness and the noisiness of the environment.

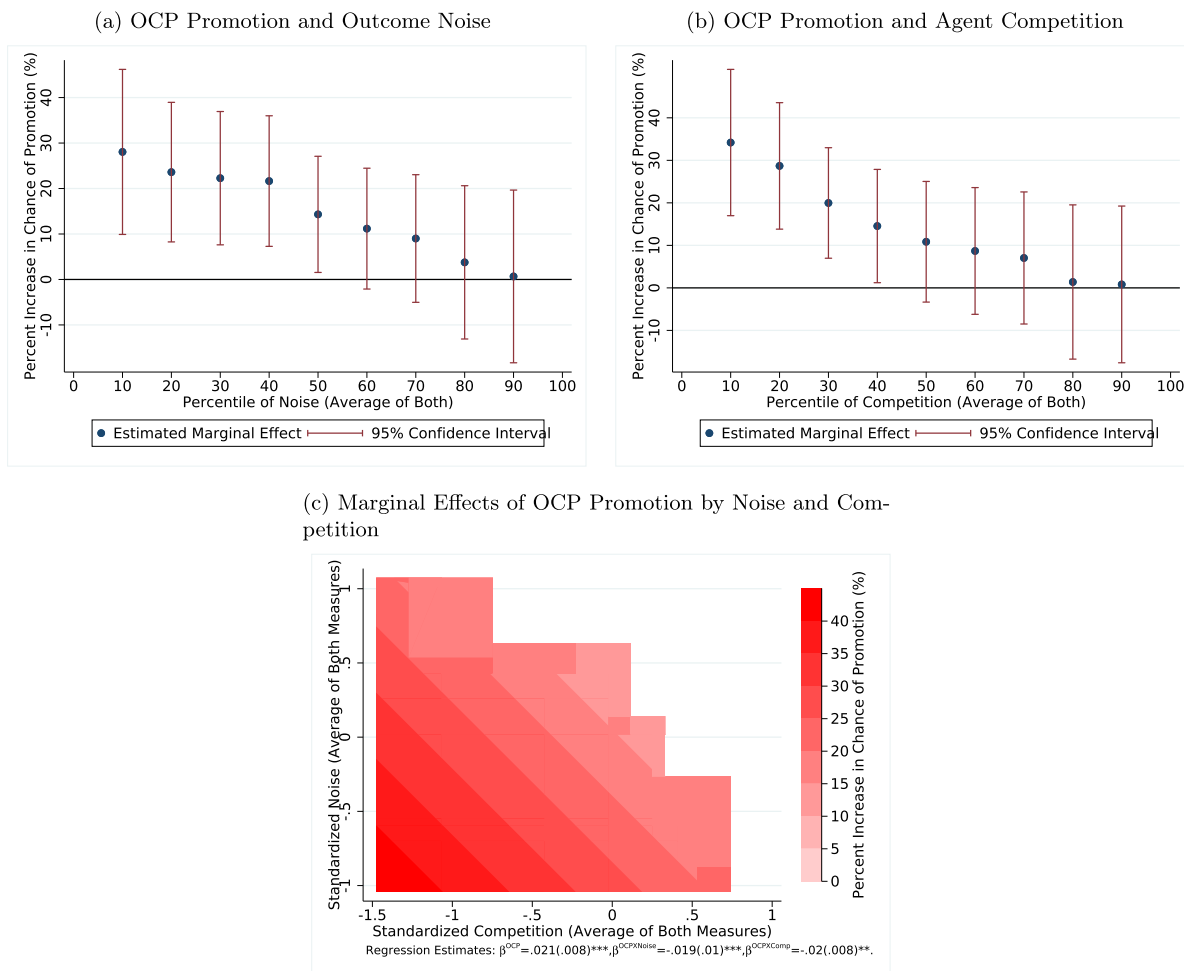


Fig. 6. Heterogeneous effects of OCP on promotion.

Notes: In Fig. 6a, the x-axis represents different percentiles of the distribution of the average noisiness measure; the higher the percentile, the noisier the signal. The y-axis is the predicted percentage change in the probability of promotion of increasing OCP performance by 1 per 1000. The figure is plotted using estimates in column (4) of Table 3. In Fig. 6b, the x-axis represents different percentiles of the distribution of the average competitiveness measure; the higher the percentile, the more competitive the tournament. The figure is plotted using estimates in column (4) of Table 4. Fig. 6c plots marginal effects of increasing OCP performance by 1 per 1000 on the probability of promotion using estimates from Table A.3. For ease of interpretation, the figure only displays effects that are statistically significant at the 10% level.

13 provinces saw a decrease in the target in the 1985 plan, while 21 provinces saw a decrease in the 1995 plan. The exclusion restriction is that target changes that occur as part of the five-year plans are not otherwise correlated with unobserved mayor characteristics that also affect promotion.

Table 5 presents results from this strategy.⁴² Column (1) presents the first-stage estimate, which displays a strong and positive correlation between decreases in the targets and reported OCP performance; the F-value of the first-stage coefficient is 21. Table A.8 shows that the instrument is not correlated with log GDP and shows that the first stage is robust across different specifications, including ones that do not control for log GDP. We also find in Table A.9 that the instrument is uncorrelated with other measures of economic performance, including GDP growth rate, electricity usage and railway cargo volume. Column (3) presents the 2SLS estimate, which is slightly larger than the OLS estimate in column (2). Decreasing the rate of natural increase by 1 per 1000 increases the chance of promotion by 2.8 percentage points,

which represents a 21% increase in the probability of promotion. This estimate has a similar economic magnitude, falls within the range of estimates in the previous section, and is not statistically different from the OLS estimate. However, one interpretation of a larger effect is that well-connected candidates for promotion might be assigned to “problem places” with larger challenges. An alternative interpretation is that the 2SLS estimate might reduce measurement error in OCP performance, leading to a larger estimate. As is often the case with IV estimates, these interpretations are at best speculative given the loss in statistical precision.

While the exclusion restriction is fundamentally untestable, we provide corroborating evidence of its plausibility. First, one concern is that province targets are set to favor a particular mayor within a province. While this does not seem feasible given that most provinces have five or more prefectures, we nevertheless analyze the potential for this concern. For example, if politically-connected mayors performed worse in birth control than unconnected ones, the target might be raised to help connected mayors improve their performance. Although an increase in target could favor every mayor within a province, only those staying in office after the change experience the benefit. We test this possibility in Table A.10, Panel A. We split the sample into two parts: column

⁴² Note that the sample is smaller than that in column (4) of Table 2 because the target data are unavailable in a few years in a few provinces.

Table 4
Heterogeneous effects of reported OCP performance on promotion by competitiveness.

	(1)	(2)	(3)	(4)
OCP Performance Reported	0.017* (0.010)	0.018** (0.008)	0.018** (0.009)	0.019** (0.008)
OCP Performance Reported X Competitiveness (Tenure)		-0.012 (0.007)		
OCP Performance Reported X Competitiveness (1-Pr Rate)			-0.018*** (0.006)	
OCP Performance Reported X Competitiveness (Average)				-0.022*** (0.008)
Log GDP	0.172*** (0.056)	0.175*** (0.054)	0.171*** (0.054)	0.174*** (0.056)
Observations	1483	1483	1483	1483
R ²	0.584	0.586	0.589	0.589
Mean of Dependent Variable	0.132	0.132	0.132	0.132
Year FE	Y	Y	Y	Y
Person FE	Y	Y	Y	Y
City FE	Y	Y	Y	Y
Prefecture-year controls	Y	Y	Y	Y
Marginal Effect at 25th pctl		.023*** (.008)	.034*** (.011)	.038*** (.01)
Marginal Effect at 50th pctl		.02*** (.008)	.015* (.009)	.014 (.009)
Marginal Effect at 50th pctl (SE)				
Marginal Effect at 75th pctl		.008 (.012)	.007 (.01)	.003 (.011)
Marginal Effect at 75th pctl (SE)				
One-Sided Test of Interaction (p-val)		0.057	0.003	0.004

Notes: OCP performance reported is the target for the rate of natural increase minus the reported rate of natural increase. Competitiveness (Tenure) is measured by the average tenure of upper-level officials at the province level. Competitiveness (1-Pr Rate) is (1-promotion rate by province). Both competitiveness measures are standardized. Competitiveness (Both) is the average of Competitiveness (Tenure) and Competitiveness (1-Pr Rate). Tenure fixed effects and migration controls are included in all regressions. Prefecture-year controls include log(population), percentage of urban population and log(government investment). Standard errors are clustered at the province level and are bootstrapped. * significant at 10% level; ** significant at 5% level; *** significant at 1% level.

Table 5
Instrumental variable regression of promotion on reported OCP performance.

	OCP Performance Reported	Promotion = 1	
	(1) First Stage	(2) OLS	(3) 2SLS
Change in Birth Rate Target	0.563*** (0.124)		
OCP Performance Reported		0.017* (0.010)	0.028* (0.017)
Log GDP	0.279 (0.511)	0.183*** (0.056)	0.179*** (0.051)
Observations	1519	1371	1519
R ²	0.80	0.47	0.58
First-Stage F-Stat			20.65
Year FE	Y	Y	Y
Person FE	Y	Y	Y
City FE	Y	Y	Y
Prefecture-year controls	Y	Y	Y

Notes: OCP performance reported is the target for the rate of natural increase minus the reported rate of natural increase. Tenure fixed effects and migration controls are included in all regressions. Prefecture-year controls include log(population), percentage of urban population and log(government investment). Standard errors are clustered at the province level and are bootstrapped. Table A.10 tests potential violations of the exclusion restriction. Tables A.11 and A.12 explore heterogeneous effects by noisiness and competitiveness using this IV strategy. * significant at 10% level; ** significant at 5% level; *** significant at 1% level.

(1) uses the sub-sample in provinces and years with an increase in the target, and column (2) uses the sub-sample in provinces and years with a decrease in the target. In both sub-samples, we fail to find evidence that the OCP performance prior to a target change is correlated with whether they stay in office after the change. A second concern is that if connected mayors anticipate a change in the target, there could be

selection on whether they are promoted prior to the change. We test this hypothesis in Panel B and we do not find a statistically significant correlation between future changes in the targets and promotion, suggesting that the target is not changed to favor some (connected) mayors. Finally, one might be concerned that mayors respond to a decrease in target by reporting lower rates of natural increase. Panel C indicates

that changes in the target do not significantly change the difference between the rate from census data and the reported rate, which rules out this last concern.

We also test whether 2SLS estimates of the effect of OCP performance on promotion differ by noisiness and competitiveness. In Table A.11, we include the interaction of OCP performance with our noisiness measures, which we instrument with interacted versions of decreases in targets. Consistent with the results in Table 3, the OCP performance has a smaller effect on promotion in noisier environments. In Table A.12, we estimate the average effect of OCP performance on promotion allowing for heterogeneous effects across regions with low versus high competitiveness.⁴³ We instrument for low competition interacted with OCP performance using low competition interacted with the decrease in target, and for high competition interacted with OCP performance using high competition interacted with the decrease in target. The 2SLS results are consistent with the baseline results that the promotion incentive is smaller in more competitive environments.

4.2.3. Using estimated prefecture-level targets

To the best of our knowledge, mayors are evaluated using province-level targets. However, a potential concern is that the relevant target for mayors' evaluations is subcontracted to prefecture-level governments. If this were the case, measurement error in OCP performance using province-level targets might affect our estimates. We address this concern by estimating an allocation rule of targets across provinces and using this rule to predict targets at the prefecture level. Table A.13 reports the estimated allocation rule at the province level. We find that higher targets are allocated to provinces with a higher number of women of reproductive age (15–45) and a higher fraction of rural women. Using these estimates and the same demographic measures at the prefecture level, we predict a prefecture-level target.

Table A.14 shows that our main results are robust to using our estimated province-level targets. Column (1) shows that controlling for the prefecture-level demographic variables used in Table A.13 results in similar average estimates to those in Table 2. Columns (2)–(4) consider different linear combinations of the province- and prefecture-level measures, which result in statistically and economically similar estimates. Further, columns (5)–(7) use the average of these measures in the comparative static analysis, which result in similar interactions with our measures of noise and competitiveness, confirming the role of meritocracy in the promotion rule.

Another potential concern is that some provinces are dominated by a few mayors and that the targets are set in order to favor them. To address this concern, we create two measures of mayor dominance: (1) the share of a province's population represented by a given mayor and (2) the Herfindahl-Hirschman Index (HHI) for each province. On average, mayors represent 16% of a province's population. Similarly, the average HHI is 21%, which implies that, on average, provinces are divided into 5 equally-populated areas. In Table A.15, we show that excluding mayors that account for more than 25% of a province's population or relatively more concentrated provinces (those with HHI > 33%) does not affect the results of Table 2. These results show that most provinces are not dominated by a few or even a single mayor. Therefore, it is unlikely that the concerns of potentially dominant mayors could sway the setting of targets.

4.2.4. Measurement error

We now perform a set of analyses where we leverage our two measures of targets (province- and prefecture-level) to assuage concerns of

potential measurement error. Following the repeated measures literature (see, e.g., Bound et al., 2001), we use one measure as an instrument for the other. Tables A.16 and A.17 present results from both iterations of this procedure. In both cases, we find very similar estimates of the main effect of OCP performance on the probability of promotion as well as similar comparative static results that support our conclusion that $\beta_1 > 0$, $\beta_2 < 0$, and $\beta_3 < 0$.

4.2.5. Alternative specifications

We also explore whether lagged OCP performances in the past four years affect our contemporaneous estimate in Table A.18 and find that our main result is robust to including lagged OCP performance in the regression. Additionally, Table A.3 shows that our results are robust to using different levels of clustering, Table A.5 shows that our results are robust to including interactions between OCP performance and log GDP, and Table A.4 shows that we obtain very similar results when we include province-by-5-year-plan fixed effects and province-by-year fixed effects.

5. Empirical evidence of manipulation and screening efficacy

The previous section shows that reported OCP performance has a positive and substantively large effect on the probability that a given mayor is promoted. Moreover, tests of comparative statics further support the view that the promotion rule is consistent with a meritocratic objective. However, it is still an open question whether such a tournament mechanism with non-contractible output is able to screen successfully for high-ability mayors, and whether the population audits have an effect on misreporting behavior.

5.1. Effects of audits on output and data manipulation

In this section we analyze the misreporting behavior of mayors. As we discuss in Section 3, we find significant disparities between the reported OCP performance and the Census OCP. While this is consistent with strategic misreporting, it could also be the case that this difference is influenced by measurement error. In order to show that mayors engage in strategic misreporting, we focus on how misreporting changes in years when mayors have a lower incentive to misreport. As discussed in Section 1.3, the central government uses population census and national fertility surveys to investigate the actual rates and the credibility of reported rates, which are organized at the province level. The audit year is the year before the census or fertility survey when the actual rates are fully observed. Equation (5) tests whether the difference between reported rates and the rates using the census data is smaller one year prior to the census or national fertility survey (i.e. the audit year):

$$\text{Birth Rate}_{cpt}^{\text{census}} - \text{Birth Rate}_{cpt}^{\text{reported}} = \delta \text{Audit}_t + f(t) + X_{icpt}\gamma + \mu_i + \eta_{cp} + \varepsilon_{icpt}, \tag{5}$$

where the binary variable Audit_t is equal to 1 in years 1987, 1989, 1991, 1994, and 1996. We include a flexible year trend $f(t)$. If the difference indeed suggests data manipulation, we should observe that $\delta < 0$. If this is the case, we further examine whether the decrease in the difference from these two data sources comes from higher reported rates, suggesting less manipulation, or from lower actual rates in census, indicating actual improvement in OCP enforcement.

Table 6 presents estimates of Equation (5). In column (1), we find that the difference between reported rates and actual rates from the census is smaller in audit years. This is consistent with the view that birth rates are manipulated in non-audit years. We further examine whether the decrease in the difference from these two data sources comes from higher reported rates, suggesting less manipulation, or lower actual rates, indicating actual improvement in OCP enforcement. Columns (2)

⁴³ We define low competitiveness as the lower tercile of the distribution.

Table 6
Effects of population audits.

	(1) Census - Reported Birth Rates	(2) Reported Birth Rate	(3) Census Birth Rate	(4) Log GDP
Audit Year	-0.325** (0.164)	0.039 (0.132)	-0.286* (0.164)	0.174*** (0.014)
Observations	1524	1524	1524	1524
R ²	0.84	0.81	0.92	0.98
Cubic Year Trend	Y	Y	Y	Y
Person FE	Y	Y	Y	Y
City FE	Y	Y	Y	Y
Prefecture-year controls	Y	Y	Y	Y

Notes: Misreporting is measured as the difference: Census - Reported Birth Rates. Audit years include the year before the census year in 1990 and 1995, and the year before the national fertility survey in 1988, 1992, and 1997. Tenure fixed effects and migration controls are included in all regressions. Prefecture-year controls include log(population), percentage of urban population and log(government investment). Standard errors are clustered at the province level and are bootstrapped. * significant at 10% level; ** significant at 5% level; *** significant at 1% level.

and (3) suggest that most of the decrease is attributable to mayors exerting more effort to lower the actual rate. Finally, if lower birth rates and higher GDP growth rate are substitutes, we would expect that, in audit years when mayors reported higher rates of natural increase (thus doing worse on the OCP dimension), mayors might work harder to improve GDP or report higher GDP numbers. We find that this is indeed the case in column (4).⁴⁴

We can use the estimates in Tables 2 and 6 to provide an assessment of the relative costs of increasing monitoring. First, we compute the government’s perceived benefit from reducing the birth rate. Given our estimates from Table 2, the government is willing to trade-off 10% in GDP growth for a 1 in 1000 birth rate reduction. Second, we compute the benefits of an additional audit. Following Table 6, we find that an audit resulted in -0.29 fewer births per 1000 people. This implies that the government’s willingness to pay for an audit is approximately 2.9% of local GDP. We view this back-of-the-envelope calculation as suggestive of two insights. First, given the high willingness to pay for an audit, this implies that the government faces significant costs from not being able to perfectly verify OCP performance. Second, given that the government only conducted audits in five of the fifteen years in our dataset, we infer that conducting additional audits (to increase the degree of verifiability) was simply too expensive for the government.

5.2. Screening ability in the presence of misreporting

Our model predicts that when mayors manipulate reported birth rates, the expected ability of the promoted mayor will be lower than when output is contractible. In the extreme case where output is completely non-verifiable, the ability of the promoted mayor will approach the population average. We test this hypothesis in Equation (6):

$$Promoted_{i_{cpt}} = \beta_4 OCP_{cpt}^{census} + X_{i_{cpt}}\gamma + \mu_i + \eta_{cp} + \lambda_t + \varepsilon_{i_{cpt}}, \tag{6}$$

where the key regressor of interest is the OCP performance measure from census data, measured as $OCP_{cpt}^{census} = Target_{pt} - BirthRate_{cpt}^{census}$. In the case where output is non-verifiable, our theory predicts that $\beta_4 \approx 0$.

Table 2 presents estimates of β_4 , where we control for the same controls as Equation (2) as well as for $OCP_{cpt}^{reported}$. From column (1) through column (6), results from all specifications suggest that actual OCP performance is not significantly predictive of promotion. These findings imply that promoted mayors are not significantly more able to

lower the actual rate than mayors who are not promoted. This shows that meritocracies may not screen high ability agents when output is imperfectly verifiable. In addition to showing that $\beta_4 \approx 0$, we also test whether the effect of $OCP_{cpt}^{reported}$ is statistically larger than the effect of OCP_{cpt}^{census} , that is, that $\beta_1 > \beta_4$. We report p-values from this test in the bottom of Table 2. The hypothesis that $\beta_1 < \beta_4$ is rejected with p-values between 0.06 and 0.08.

Table 2 uses the best information available in the census to account for migration. In Table A.19, we explore the potential for migration to explain our null results by studying the effect of removing migration controls.⁴⁵ Similarly, in Table A.20 we follow the robustness checks described in Section 4.2. In both cases, we find that migration controls, controlling for alternative measures of economic performance, 2SLS estimates, prefecture-level targets, and corrections for measurement error deliver economically small and statistically insignificant correlations between promotion and actual performance. This suggests that the lack of correlation between actual OCP performance and promotion is not a statistical anomaly, but rather a result of the impaired capacity of the tournament mechanism to screen for ability when output is non-contractible and reported performance is manipulated.

6. Conclusion

This paper analyzes the role of meritocracy in determining the promotion of mayors in China. We document that, despite potential for corruption through political connections, promotion rules are partly driven by perceived performance in implementing the OCP. Moreover, we show that the relationship between performance and promotion is determined by a desire to screen high-ability mayors for higher office. While we confirm that observed promotion decisions are consistent with a meritocratic objective, the efficacy of this screening mechanism is weakened by mayors’ ability to manipulate reported outcomes. Empirically, we find that mayors manipulate less in audit years (that is, when monitoring is increased), which is consistent with the importance of OCP as a performance metric. Nonetheless, we find that audits are not able to resolve the fundamental problem of non-contractible output.

The combination of theory and empirical analysis makes our findings particularly compelling and demonstrates the importance of inter-

⁴⁴ Note that the misreporting of mayors above the target does not respond to audits, as we show in Table A.23.

⁴⁵ This analysis investigates the potential measurement error from not directly observing migration before 1995 (when the migration rate was below 4%). We repeat the specification in column (4) of Table 2, with various incomplete controls for migration and using the same sub-sample, and report the results in column (1) through column (3) in Table A.19; the main estimate is not sensitive to including these controls.

preting empirical results through the lens of a rigorous model of incentives. Without guidance from our model, the applied econometrician could arrive at the mistaken conclusion that meritocracy was not a driving force in the Chinese government, as promoted mayors do not appear to be of higher ability than mayors who are not promoted. However, by testing more subtle predictions of our model, we are able to separate the desire to implement the OCP from the meritocratic objective, a distinction which previous studies of promotion based on other measures of performance were not able to address.

While our study focuses foremost on the implementation of the OCP, it also yields interesting lessons for the literature studying the relation between economic growth and cadre promotion. Our model suggests that promotions are less likely to stem from a meritocratic motive in cases where the government places a relatively higher value on output production. For instance, it may be the case that economic growth is valued much more as a production objective than as a screening mechanism. From the perspective of the OCP, a similar logic implies that the government may place a lower importance on this policy than was previously thought. Instead, the enduring permanence of the OCP may be due to the fact that the government believed this task helped improve the screening of mayors. This discussion highlights two facts that are evinced by our study. First, identifying the role of meritocracy in selection requires an understanding of the trade-offs faced by the principal across different incentive forms. Second, the reason why performance in policy implementation is linked to promotion is important for understanding the provenance of those policies as well as the potential for their reform.

We conclude by noting that, while critics of the implementation of the OCP point toward local government promotion incentives as a cause for human rights abuses, including forced abortions and sterilizations (see, e.g., Wong, 2012), the alternative piece rate compensation mechanism would likely lead to more such cases as effort from lower-ability mayors would likely increase. Moreover, one consequence of the non-contractibility and manipulation of the rates of natural increase is that marginal incentives may have little or no effect on actual birth rates, even though the policy on the whole may lead to human rights abuses.

Appendix A. Supplementary data

Supplementary data related to this article can be found at <https://doi.org/10.1016/j.jdeveco.2019.06.003>.

References

- Acemoglu, D., 2005. Politics and economics in weak and strong states. *J. Monet. Econ.* 52, 1199–1226.
- Baker, G., Gibbs, M., Holmstrom, B., 1994. The internal economics of the firm: evidence from personnel data. *Q. J. Econ.* 109 (4), 881–919.
- Banister, J., 1992. China: Recent Mortality Levels and Trends. Paper Presented at the Annual Meeting of the Population Association of America, 30 April–2 May, Denver, Colorado.
- Bell, D., 2015. *The China Model: Political Meritocracy and the Limits of Democracy*. Princeton University Press.
- Bell, D., Li, E., 2012 November 11. In Defence of How China Picks its Leaders. *The Financial Times*.
- Besley, T.J., Persson, T., 2010. State capacity, conflict and development. *Econometrica* 78 (134).
- Birney, M., 2014. Decentralization and veiled corruption under China's rule of mandates. *World Dev.* 53, 55–67.
- Bobonis, G., Fuertes, L.R.C., Schwabe, R., 2013. Monitoring Corruptible Politicians. *American Economic Review*. forthcoming.
- Bound, J., Brown, C., Mathiowetz, N., 2001. Measurement error in survey data. In *The Handbook of Econometrics*.
- Chiappori, P., Salanie, B., 2000. Testing for asymmetric information in insurance markets. *J. Political Econ.* 108, 56–78.
- Choukhmane, T., Courdacier, N., Jin, K., 2014. The One Child Policy and Household Saving. mimeo.
- Cullen, J.B., Reback, R., 2006. Tinkering toward accolades: school gaming under a performance accountability system. In: Gronberg, T., Jansen, D. (Eds.), *Improving School Accountability: Check-Ups or Choice*, Advances in Applied Microeconomics, vol. 14. Elsevier Science, Amsterdam, pp. 1–34.
- Ebenstein, A., 2010. The missing girls of China and the unintended consequences of the one child policy. *J. Hum. Resour.* 45 (1), 87–115.
- Ferraz, C., Finan, F., 2008. Exposing corrupt politicians: the effects of Brazil's publicly released audits on electoral outcomes. *Q. J. Econ.* 123 (2), 703–745.
- Ferraz, C., Finan, F., 2011. Electoral accountability and corruption: evidence from the audits of local governments. *Am. Econ. Rev.* 101 (4), 1274–1311.
- Fisman, R., Wang, Y., 2015. *The Economics of Death Ceilings*. Working Paper.
- Hardee-Cleveland, K., Banister, J., 1988. Fertility policy and implementation in China, 1986–88. *Popul. Dev. Rev.* 14 (2), 245–286.
- Holmstrom, B., 1999. Managerial incentive problems: a dynamic perspective. *Rev. Econ. Stud.* 66 (1), 169–182.
- Holmstrom, B., Milgrom, P., 1991. Multitask principal-agent analyses: incentive contracts, asset ownership, and job design. *J. Law Econ. Organ.* 7, 24–52.
- Jacob, B., Levitt, S., 2003. Rotten apples: an investigation of the prevalence and predictors of teacher cheating. *Q. J. Econ.* 118, 843–877.
- Jia, R.-X., 2014. Pollution for Promotion. mimeo.
- Jia, R.-X., Kudamatsu, M., Seim, D., 2015. Political selection in China: complementary role of connections and performance. *J. Eur. Econ. Assoc.* 13 (4), 631–668.
- Johansson, S., Nygren, O., 1991. The missing girls of China: a new demographic account. *Popul. Dev. Rev.* 17 (1), 35–51.
- Landry, P., 2008. *Decentralized Authoritarianism in China: the Communist Party's Control of Local Elites in the Post-Mao Era*. Cambridge University Press, Cambridge and New York.
- Lazear, E., Rosen, S., 1981. Rank-order tournaments as optimum labor contracts. *J. Political Econ.* 89 (5), 841–864.
- Li, H.-B., Zhou, L.-A., 2005. Political turnover and economic performance: the incentive role of personnel control in China. *J. Public Econ.* 89 (910), 1743–1762.
- Maskin, E., Qian, Y.-Y., Xu, C.-G., 2000. Incentives, information, and organizational form. *Rev. Econ. Stud.* 67 (2), 359–378.
- Olken, B., 2006. Monitoring corruptible politicians: evidence from a field experiment in Indonesia. *J. Political Econ.* 115 (2), 200–249.
- Persson, P., Zhuravskaya, E., 2015. The limits of career concerns in federalism: evidence from China. *Journal of the European Economic Association*. forthcoming.
- Prendergast, C., 1999. The provision of incentives in firms. *J. Econ. Lit.* 37, 7–63.
- Qian, N., 2009 May. Quantity-quality and the one child policy: the only-child disadvantage in school enrollment in rural China. Working Paper 14973. National Bureau of Economic Research.
- Rasul, I., Rogger, D., 2016 January. Management of bureaucrats and public service delivery: evidence from the Nigerian civil service. Technical Report No. DP11078., CEPR Discussion Paper.
- Rochet, J.-C., Stole, L., 2003. The economics of multidimensional screening. In: Dewatripont, M., Hansen, L. (Eds.), *Advances in Economic Theory*, 7th World Congress. Cambridge University Press.
- Scharping, T., 2003. *Birth Control in China 1949–2000: Population Policy and Demographic Development*. Routledge-Curzon, London and New York.
- Shih, V., Adolph, C., Liu, M.-X., 2012. Getting ahead in the communist party: explaining the advancement of central committee members in China. *Am. Pol. Sci. Rev.* 106, 166–187.
- Wei, S.-J., Zhang, X., 2011. The competitive saving motive: evidence from rising sex ratios and saving rates in China. *J. Political Econ.* 119 (3), 511–564.
- Wong, E., 2012 July 22. Reports of Forced Abortions Fuel Push to End Chinese Law. *The New York Times*.
- Xu, C.-G., 2011. The fundamental institutions of China's reforms and development. *J. Econ. Lit.* 49 (4), 1076–1151.
- Zeng, Y., Tu, P., Gu, B., Xu, Y., Li, B., Li, Y., 1993. Causes and implications of the recent increase in the reported sex ratio at birth in China. *Popul. Dev. Rev.* 19 (2), 283–302.
- Zheng, S., Kahn, M.E., Sun, W., Luo, D., 2014. Incentives for China's urban mayors to mitigate pollution externalities: the role of the central government and public environmentalism. *Reg. Sci. Urban Econ.* 47, 61–71. St: Tribute to John Quigley.
- Zitzewitz, E., 2012. Forensic economics. *J. Econ. Lit.* 50 (3), 731–769.