Geography, Transparency and Institutions*

Joram Mayshar† Omer Moav‡ Zvika Neeman§

July 11, 2014

Abstract

We propose a theory by which geographic variations in the transparency of the production process explain cross-regional differences in the scale of the state, in its hierarchical structure, and in property rights over land. The key linkage between geography and these institutions, we posit, is via the effect of transparency on the state's extractive capacity. We apply our theory to explain institutional differences between ancient Egypt and ancient Upper and Lower Mesopotamia. We also discuss the relevance of our theory to analyses of the deep rooted factors affecting economic development and the growth of taxation in the modern age.

KEYWORDS: *Geography, Transparency, Institutions, Land Tenure, State Capacity, State Concentration*

JEL CLASSIFICATION NUMBERS: *D02, D82, H10, O43*

---

*We are grateful for comment received from Daron Acemoglu, Bob Allen, Josh Angrist, Eddie Dekel, Diana Egerton-Warburton, Christopher Eyre, James Fenske, Oded Galor, Maitreesh Ghatack, Jeremy Greenwood, James Malcomson, Andrea Matranga, Jacob Metzer, Stelios Michalopoulos, Motty Perry, Torsten Persson, Herakles Polemarchakis, Louis Putterman, Debraj Ray, Ariel Rubinstein, Yona Rubinstein, Larry Samuelson, Matthew Spigelman, Yannay Spitzer, Nathan Sussman, Juuso Valimaki, Joachim Voth, David Weil, and from seminar participants in Royal Holloway University of London, Ben-Gurion University, The Hebrew University of Jerusalem (department of economics), The Hebrew University of Jerusalem (department of archeology) , Bilkent University, Tel-Aviv University, University of Aberdeen, University of Edinburgh, ISET, LSE/UCL growth-development seminar, University of Warwick, University of Alicante, Oxford University, Harvard/MIT applied theory seminar, NYU, Brown, Indiana, Di Tella, Copenhagen, Odense, Sant'Anna Pisa, and conference participants at The Royal Economic Society Conference, Royal Holloway, April 2011; CEPR Public Policy Symposium, Zurich, May 2011; Minerva Summer School in Economic Growth, Jerusalem, June 2011; Barcelona II Workshop "Towards Sustained Economic Growth", October 2011; NBER Summer institute, "Inequality and Growth", Cambridge Massachusetts July, 2012, Conference of the Economic History Association of Israel, Bar Ilan University, December 2012; CAGE/CEPR Conference Long-Run Growth: Unified Growth Theory and Economic History, University of Warwick, January 2013; CEPR Economic History Symposium; Perugia, April 2013 Warwick Summer School in Economic Growth, Coventry, July 2013; 57. EEA ESEM, Gothenburg, August 2013.

†Department of Economics, Hebrew University of Jerusalem. Email: msjoram@huji.ac.il.

‡Department of Economics University of Warwick, Department of Economics Hebrew University of Jerusalem, CAGE and CEPR. Email: omer.moav100@gmail.com; Moav's research is supported by the Israel Science Foundation (Grant No. 73/11).

§Eitan Berglas School of Economics, Tel-Aviv University, Email: zvika@post.tau.ac.il

1 An agency model with differential transparency

Consider a state with a given area of arable land, which is divided into plots. Each plot is allocated to one risk neutral agent-tenant. We model the principal (the state) as an absentee landlord who designs a contract that maximizes her expected periodic income, which is given by the total output produced by all the agents, less any payments made to the agents and the costs of replacing dismissed agents.[1] Each agent decides how much effort to exert. His payoff is the payment received from the principal, less his cost of effort.

We first characterize the optimal contract between the principal and a single agent in a simple basic model. We then extend the basic model to a two-tiered hierarchy with multiple agents.[2]

1.1 The basic model

We consider a Principal-Agent model that has the following characteristics. Both the output that is produced by the agent and the agent's choice of effort can be either low or high: $Y \in \{L, H\}$, and $e \in \{l, h\}$, respectively. The state of nature is also binary and can be either good or bad: $\theta \in \{G, B\}$. The annual output is a function of the effort exerted by the agent and the state of nature. We assume that it is high if and only if both the state of nature is good and the agent exerts high effort:

$$Y = \begin{cases} H & \text{if } e = h \text{ and } \theta = G; \\ L & \text{otherwise.} \end{cases}$$

The agent chooses the level of effort before he learns the state of nature.[3] The ex-ante probability that the state of nature is good is denoted by: $p \in (0, 1)$. After the agent chooses the level of effort, both the agent and the principal observe a public signal about the state of nature: $\sigma \in \{\tilde{G}, \tilde{B}\}$. The accuracy of this signal, $q \in [1/2, 1]$ is such that

$$Pr(\tilde{G}|G) = Pr(\tilde{B}|B) = q;\ Pr(\tilde{G}|B) = Pr(\tilde{B}|G) = 1 - q.$$

[1]We thus follow Olson (1993) and model the state as engaged in the expropriation of its subjects, without explicit consideration of the use of the revenue for the provision of public goods. This abstraction is not crucial for our arguments. Besides, beyond the provision of security, which we take for granted because it also serves the interests of the state, it seems to us to be a rather reasonable simplification. Besley and Ghatak (2009, p. 4560) claim that starting at the 14th century CE, "Expropriations by government are a fact of historical experience"; Ma (2011) offers a similar perception of the state in imperial China.

[2]Throughout our analysis we assume that plot size and therefore also the size of the population are given. In online Technical Appendix F, we generalize the model to include an endogenous population size and an endogenous plot size that is determined by the principal to maximize its income. The main qualitative results are unchanged.

[3]In practice, both the agent's effort and the relevant state of nature for agriculture are vectors whose components are distributed over the agricultural seasons. In online Technical Appendix C we show that if the agent learns the state of nature before exerting effort, then the payoff to the Principal is higher. No other qualitative changes to the model are implied.

The level of accuracy q represents the degree of transparency of production in our model. If $q = 1$ then the signal perfectly reveals the state of the world; if $q = 1/2$ then the signal is uninformative.

We denote the annual cost (in units of output) of providing for the agent (and his family) until the next harvest period by $m + \gamma$, where $m \geq 0$ is the cost of subsistence in case the agent exerts (costless) low effort, and $\gamma > 0$ is the annual cost of exerting high effort. We assume that output is in any case larger than the cost of providing for an agent who exerts high effort: $L \geq m + \gamma$.

The agent's only alternative mode of employment outside agriculture is work as a domestic servant. We normalize his utility in this case to zero. The agent's annual utility as a tenant farmer equals his expected income, denoted by I, less the cost of subsistence and effort. Thus, the agent's annual utility if he exerts high and low effort is given by $I - (m + \gamma)$ and $I - m$, respectively. We assume that the agent has no other sources of income or wealth, that he cannot save, and that he cannot borrow. We denote the agent's intertemporal discount factor by $\delta \in [0, 1)$.

The principal employs the following incentive scheme. If output is low, the principal pays the agent a basic wage ω. If output is high, then the principal pays the agent $\omega + b$, where $b \geq 0$ is a bonus payment. The basic wage ω has to sustain an agent who exerts effort until the next harvest, hence $\omega \geq m + \gamma$.

When output is high the principal retains the agent. The agent is also retained when output is low but the signal indicates that the state of nature is bad ($\sigma = \tilde{B}$). But if output is low and the signal indicates that the state of nature is good ($\sigma = \tilde{G}$), then the principal may dismiss the agent and replace him with another.

We denote the probability with which the agent is dismissed when output is low and the signal indicates that the state of nature was good by d. For simplicity, we assume that the principal employs a pure strategy, namely $d \in \{0, 1\}$.[4] If the agent is dismissed, then the principal incurs a fixed cost $x > 0$ that represents the cost of dismissal and the present value of lost output during the training of a new agent.

This simple incentive scheme (plus the two extensions) is the most general incentive scheme available to a principal who can only pay to and/or dismiss the agent and who can only condition its payment and dismissal probability on histories of bounded finite length if the agent has no property and therefore cannot be fined.

We refer to the bonus payment b as a 'carrot' and to the possibility of dismissal d as a 'stick.'

[4] In online Technical Appendix D we extend the model to the case where $d \in [0, 1]$. In Appendix E we consider an alternative extension, where the principal may warn the agent when he suspects him of shirking, and dismiss the agent only after an endogenously determined number of warnings. The qualitative results of the model regarding the effect of transparency q on the optimal contract are unchanged in both extensions.

Thus, the solution of the Principal-Agent problem described here strikes an optimal balance between the use of a carrot and a stick as incentive devices. The fact that the principal is restricted to set d equal to either zero or one implies that only two types of contracts may be optimal. We refer to the contract where $d = 0$ as the 'pure carrot' contract, and to the contract where $d = 1$ as the 'stick and carrot' contract. We denote this pair of contracts with subscripts c and s respectively. Under the 'pure-carrot' contract, the agent is never dismissed and is incentivized only through bonuses. Under the 'stick and carrot' contract the agent is dismissed whenever output is low but the signal is good ($Y = L$, $\sigma = \tilde{G}$).[5]

The optimal balance between the carrot and the stick depends on the transparency of production, or the precision of the public signal q, as described in the following proposition.

Proposition. If $x > p\delta\gamma / (1 - \delta/2)(1 - p)$, then the optimal contract that is selected by the principal has the following properties:

1. the agent's basic wage is set at its lowest possible value, or $\omega = m + \gamma$
2. There exists a threshold $\hat{q} \in (1/2, 1)$ such that:

   if $q < \hat{q}$, then the optimal contract is a 'pure carrot' contract:

$$d_c = 0 \text{ and } b_c = \gamma/p;$$

   if $q > \hat{q}$, then the optimal contract is a 'stick and carrot' contract:

$$d_s = 1 \text{ and } b_s = \frac{\gamma}{p}\left(1 - \frac{pq\delta}{1-\delta(p+q-2pq)}\right);$$

   if $q = \hat{q}$, then both contracts above are optimal.

If $x \leq p\delta\gamma / (1 - \delta/2)(1 - p)$, then either the 'stick and carrot' contract or dismissal of the agent upon observation of low output are optimal.

Proof. Denote by V the present value of the agent's utility from being employed in agriculture in a stationary equilibrium where he exerts high effort every period. The fact that the agent's utility upon dismissal is zero implies that:

$$V = [\omega + pb - m - \gamma] + [1 - \Pr(dismiss|e = h)]\delta V, \tag{1}$$

[5] One may argue that the principal may have an incentive to renege on the contract chosen, and to avoid paying the bonus to the agent, or to not dismiss the agent when this is called for by by the contact. This is not a concern, however, if the principal is patient and faces many agents simultanuously who are likely to believe that once the principal reneges, she will continue to do so in the future.

Denote the probability of a bad harvest and a good signal for an agent who exerts high effort by $\mu = (1-p)(1-q)$. The probability of dismissal of an agent who exerts high effort is $d\mu$. It follows from (1) that:

$$V(b,d) = \frac{\omega + pb - m - \gamma}{1 - \delta(1 - d\mu)}. \tag{2}$$

The principal's objective function (OF) is to maximize her per-period expected payoff, denoted by π,

$$\pi = \max_{b \geq 0, d \in \{0,1\}, \omega \geq \gamma,} p(H-b) + (1-p)L - \mu dx - \omega, \tag{OF}$$

subject to providing the agent with incentives to exert high effort:

$$\begin{aligned} &p[b + \delta V] + (1-p)[q + (1-q)(1-d)]\delta V + \omega - m - \gamma \geq \\ &p[q(1-d) + (1-q)]\delta V + (1-p)[q + (1-q)(1-d)]\delta V + \omega - m, \end{aligned} \tag{3}$$

where $V = V(b,d)$ as in (2). Part (1) of the proposition follows from the fact that since ω cancels out from (3) it is optimally set to $\omega = m + \gamma$. Plugging (2) into constraint (3) and simplifying yields the incentive constraint:

$$pb\left(1 + \frac{pqd\delta}{1 - \delta(1 - d\mu)}\right) \geq \gamma. \tag{IC}$$

Part (2) follows from the maximization of (OF) subject to (IC). Because the Principal sets b as low as possible, the incentive constraint is binding in the optimal solution. The threshold $\hat{q}$, is given by the unique solution in the interval $[0,1]$ of the quadratic equation that equates the values of the objective function with $d = 0$ and $d = 1$. To see that $\hat{q} > 1/2$ if $x > p\delta\gamma/(1-\delta/2)(1-p)$, rearrange this quadratic equation as:

$$\hat{q}/(1-\hat{q}) = (1-p)\,x[1 - \delta(p + \hat{q} - 2p\hat{q})]/p\delta\gamma, \tag{4}$$

and note that while the left-hand-side of (4) is increasing from zero to infinity as $\hat{q}$ increases from zero to one, the right-hand-side is positive and linear in $\hat{q}$. This implies that the intersection between the two curves is unique in the interval $[0,1]$. The condition on x is obtained by requiring that for $\hat{q} = 1/2$ the right-hand-side is larger than the left-hand-side, which is equivalent to $\hat{q} > 1/2$.

Finally, in the analysis above we only considered two pure strategies, the third pure strategy of dismissal of the agent upon observation of low output regardless of the signal is dominated by the 'pure carrot' contract if $x > \delta p\gamma/(1-p)$. Thus, it is never optimal in the range where $x > p\delta\gamma/(1-\delta/2)(1-p)$. □

The switch of the optimal contract from 'pure carrot' to 'stick and carrot' when the quality of information improves captures the essence of our claims. The logic behind it is simple. A principal relying on a 'stick' to incentivize the agent has to incur the cost x whenever a dismissal takes place. But given that the agent is incentivized to exert effort, dismissal is in fact a "wasteful mistake" that occurs with probability $\mu = (1-p)(1-q)$. The probability of dismissal, and so also the expected cost of dismissal, μx, decreases when the quality of information q improves. When dismissal is sufficiently costly to the principal, incentivizing the agent through a stick is beneficial for the principal only when q is large enough. The threshold $\hat{q}$ is determined so that it exactly balances the expected cost of dismissal μx with the expected savings to the principal due to a smaller bonus.

The effect of transparency on income and its allocation

If the economy is less transparent ($q < \hat{q}$), the principal optimally refrains from ever dismissing the agent. In this case, the contract is socially efficient and the expected income of both the principal and the agent is independent of q. Under this 'pure-carrot' regime the expected income of the agent, I_c, and the principal, π_c, are:

$$I_c = m + 2\gamma \text{ and } \pi_c = p(H-L) + L - 2\gamma - m,$$

and their combined expected income is:

$$I_c + \pi_c = p(H-L) + L.$$

In contrast, if the economy is sufficiently transparent ($q > \hat{q}$), then the optimal contract is a 'stick and carrot:'

$$I_s = m + 2\gamma - \frac{pq\delta\gamma}{1-\delta(p+q-2pq)},$$

$$\pi_s = p(H-L) + L - m - 2\gamma + \frac{pq\delta\gamma}{1-\delta(p+q-2pq)} - \mu x,$$

and

$$I_s + \pi_s = p(H-L) + L - \mu x.$$

The expected total income reveals that the 'stick and carrot' contract is socially inefficient because the agent is sometimes dismissed even though he works diligently. This efficiency loss, namely the expected cost of dismissal μx, declines as accuracy improves. In the limit, when the signal is accurate ($q = 1$), then the 'stick and carrot' regime becomes socially efficient.

The principal's payoff is continuous at the threshold of transparency $\hat{q}$ and increases with q thereafter. The gains to the principal from a rise in q above $\hat{q}$ are derived both from a rise in total income and from a decline in the agent's income. Indeed, it is the agent who bears the entire burden of the 'stick and carrot' regime: at the threshold accuracy, $\hat{q}$, his expected income I drops discreetly by the expected cost of dismissal: $(1-p)(1-\hat{q})x$. Past that threshold, his expected per-period income continues to decline with q. Within this range, the benefit that the agent obtains due to the reduced probability of dismissal enables the principal to reduce the bonus payment b, while still maintaining the incentive constraint. These features are summarized by Figure 1 below. The principal's expected income π as a function of accuracy q is depicted by the lower solid line. Total expected income $I+\pi$ is depicted by the upper solid line; and the difference between these two lines represents the agent's expected income.

Figure 1: Periodic expected income as a function of signal accuracy

Figure 1 adopts a simple illustrative calibration. We set: $H = 1.1$, $L = 0.6$ and $p = 0.8$, so that a bad harvest with a significantly lower crop occurs once in about every five years, and so that the expected crop size of each plot is set to one: $E(Y) = pH + (1-p)L = 1$.[6] To be consistent with tenants' output share of about two thirds and with the relative high cost of maintaining a family throughout the year, we set the subsistence cost to $m = 0.5$ and the effort cost to $\gamma = 0.1$, and thus the basic wage is $\omega = 0.6$. Given an interest rate (in grain) of one third or more in the ancient world, we set $\delta = 0.75$. Finally, we set $x = 2$, so that the present value cost of dismissing and replacing an agent is two expected crops.[7]

[6] One should think of this unit as representing about 1.5 tons of grain of output, net of the grain that is needed for seed (typically assumed to be about 15 percent of the crop) and also net of expected spoilage in storage (typically assumed to be another 10-20 percent). For a more elaborate attempt to calibrate early Near Eastern farming see Hunt (1987).

[7] With these parameters $\hat{q} > 1/2$ is achieved already with $x = 0.48$, however, in the version of the model in which d is continuous, (online technical appendix D), a higher x is required for obtaining a range of $q > 1/2$ in which $d = 0$ is optimal. Thus, for consistency, we set $x = 2$. A number of additional elements also serve to render the stick less attractive and help guarantee that $\hat{q} > 1/2$ with a much lower value of x. These include: a cost to obtain a signal and effort exerted in land maintenance (since the threat of eviction might reduce this investment, as shown empirically

It is instructive to compare the outcome when the signal fully reveales the state of the world ($q = 1$) with the outcome when the signal is highly inaccurate ($q < \hat{q}$). In both cases the diligent agent is never dismissed and the economy is efficient (As seen in Figure 1). However, the distribution of income is quite different. The agent's (gross) income falls from $I_c = m + 2\gamma$ in the range of the opaque signal to $I_s = m + 2\gamma - p\delta\gamma/[1 - \delta(1 - p)]$ when $q = 1$, since the bonus that the agent requires in order not to shirk is reduced to a minimum. The agent's utility from being employed in agriculture, namely his income net of effort, is entirely dissipated in this case if he is very patient ($\delta = 1$).

Discussion

Let us summarize our findings thus far. If the level of transparency is sufficiently low, then the agent-tenant is in a 'pure carrot' regime in which he is never dismissed. In such a regime the agent may be considered a de-facto owner of the land that he cultivates. While he may not be able to sell the land, he may nevertheless be able to transfer the right to cultivate it to another. This effective ownership is not due to any explicit legal rights or to the benevolence of the principal (the state), nor to any impediments that prevent dismissal. Rather, the agent's rights to the land stem from the fact that, given low transparency, it is optimal for the principal to refrain from dismissal.

In contrast, when the level of transparency is sufficiently high the agent-tenant is in a 'stick and carrot' regime. In that regime, when low output is observed, the farmer may be subject to dismissal, and thus cannot be considered to have ownership rights to the land that he cultivates. In this range, the principal relies more on the 'stick' and less on the 'carrot' as transparency increases. That is, under the optimal 'carrot and stick' contract, the higher the accuracy of the signal that the principal obtains, the stronger is the incentive provided by the threat of dismissal, since the probability of unjust dismissal of an agent who exerts high effort is then smaller, as is the probability of retaining a tenant who exerts low effort. In this range of sufficiently high transparency, the tenant's share of output declines with transparency, and the state's revenue increases.

These findings, and in particular the general effect of transparency on the optimal combination of the 'stick' and 'carrot', are robust and do not depend on our specific modeling assumptions. The credible threat of using the 'stick' reduces the cost of incentivizing the agent with the 'carrot.' However, to maintain credibility of the threat, punishment must be used (even if unjustifiably) whenever output is low and the signal is good. Since the probability of such punishment declines

by Deininger and Jin, 2006).

with transparency, it follows that the expected cost of including a 'stick' in the contract also declines with transparency.

The assumptions that are crucial for these results are standard in the literature. They include punishment through the threat of eviction and limited liability of the agent in the form of a lower bound on material remuneration. The specific assumption that the minimal remuneration, ω, is equal or larger than the cost of subsistence when exerting high effort, $m + \gamma$, is made here in order to simplify the exposition. A strictly positive ω is crucial for the effectiveness of the stick.

Our assumption that punishment takes the form of dismissal is crucial for the application of our model to the study of institutions. Our main departure from the existing literature is in including a signal that the principal observes and upon which she conditions the contract. This allows us to perform comparative statics with respect to the accuracy of this signal and obtain, thereby, new insights on the link between geography and institutions.

1.2 A Two-Level Hierarchy Model

We now extend our model to include two tiers of government. Extension of the model further to n tiers of hierarchy along the same lines is straightforward. For the relations between the governor and the farmers in the village under her control we keep the basic model. For the relations between the upper echelon (the king) and lower level of hierarchy (the village governor), we employ a variant of our basic model where the governor may hide output rather than exert low effort, because this seems more consistent to us with the historical record.

We attach a subscript of 1 or 2 to variables at each level of the hierarchy, from the bottom up. We assume that there are two independent state variables that determine the state of nature in each plot of land: $\theta_1 \in \{G, B\}$ is plot specific, and $\theta_2 \in \{G, B\}$ is village specific. The plot-specific states are assumed to be independent across plots, conditional on the village's specific state, and the village specific states are assumed to be independent across villages. We denote by $p_1 \in (0, 1)$ the probability that each plot of land is in a plot-specific good state, and by $p_2 \in (0, 1)$ the corresponding probability for the entire village.

As in the basic model, output in each plot can be either low or high: $Y_1 \in \{L_1, H_1\}$ and the agent's effort can be either low or high: $e \in \{l, h\}$. Plot output is assumed to be high if and only if the agent exerts high effort and both the plot's and village's states of nature are good ($\theta_1 = \theta_2 = G$). Thus, the state of nature in a specific plot is good with probability $p_1 p_2$, otherwise it is bad.

We assume that the village specific state of nature, θ_2, is revealed to both the farmer and the governor after the farmer's effort decision is made. In addition, if the village specific state is good

($\theta_2 = G$), then the governor receives plot-specific signals σ_1 for each plot in the village. These signals are accurate with probability $q_1 \in [1/2, 1]$ and are (conditionally) independent across plots.

At the higher level of the hierarchy, between the village governor and the king, we assume an analogous information structure. The king does not know the specific states θ_1 of individual plots, nor the states θ_2 for any of the villages. But he receives an independent signal σ_2 about each of the latter, whose accuracy is denoted by the probability $q_2 \in [1/2, 1]$.

The contract selected by the governor will have the same structure as in the basic model. It specifies a basic wage $\omega_1 = \gamma$, a bonus b_1 if output is high, and a dismissal probability $d_1 \in \{0, 1\}$ at a cost of x_1 to the governor, if output is low ($Y_1 = L_1$) but both the village's state and the plot's signal are good ($\theta_2 = G, \sigma_1 = \tilde{G}_1$). Thus, subject to the farmer exerting effort, he is dismissed with probability: $\Pr(dismiss|e = h) = (1 - p_1)\, p_2\, (1 - q_1)\, d_1$.

The governor's maximization problem is a variant of the principal's problem in the basic model, in which $p_1 p_2$ substitutes for p as the probability of high output, and in which the probability of dismissal is $p_2\,(1 - p_1)\,(1 - q_1)\,d_1$ instead of $(1 - p)(1 - q)d$. Thus, the governor chooses a 'pure carrot' contract ($d_{1c} = 0$) if transparency is below some threshold, $q_1 < \hat{q}_1$, and a 'stick and carrot' contract if $q_1 > \hat{q}_1$. Above $\hat{q}_1$, the expected income of the governor is increasing with q_1.[8]

We now turn to study the king's problem. We assume that the number of plots in each village is sufficiently large so that the total revenue obtained by the village governor can be substituted by their expected value. The governor's revenue is then limited to two possible outcomes, depending on the village-specific state of nature θ_2. We denote by L_2 and H_2 the governor's income in a bad year ($\theta_2 = B$) and a good year ($\theta_2 = G$) respectively. If N_1 is the number of plots in a village, then:

$$\begin{aligned} L_2 &= N_1 \left[L_1 - \gamma) \right], \\ H_2 &= H_2(q_1) = N_1 \left[p_1(H_1 - L_1 - b_1) + L_1 - (1 - p_1)(1 - q_1) d_1 x_1 - \gamma) \right]. \end{aligned}$$

The parameters b_1 and d_1 are those selected by the governor (as a function of q_1). Beyond the threshold $\hat{q}_1$, the good-year revenue H_2 is increasing in q_1.

Recall that the king receives a signal σ_2 regarding the village state θ_2, whose accuracy is denoted by q_2. As mentioned above, we assume that the governor may under-report the revenue collected to the king. That is, she may report collecting L_2, even though she in fact collected H_2. The king is assumed to employ an analogous two-edged incentive scheme to the one above: a bonus b_2 if the

[8]The corresponding bonus payments are: $b_{1c} = \gamma/p_1 p_2$ under 'pure carrot' and $b_{1s} = (\gamma/p_1 p_2)\left[1 - p_1 p_2 q_1 \delta_1 / \left(1 - \delta_1(1 - p_2) - \delta_1 p_2 (p_1 + q_1 - 2p_1 q_1)\right)\right]$ under 'stick and carrot'. If $p_2 = 1$, this is identical to the analogous expressions under the basic model.

governor reports collecting H_2, and a threat of dismissal at a cost of x_2 to the king, if the governor's report is L_2, but the signal σ_2 indicates that a large village harvest was to be expected.

The king maximizes:

$$\pi_2 = \underset{b_2 \geq 0, d_2 \in \{0,1\}}{max} p_2(H_2 - b_2) + (1 - p_2)[L_2 - (1 - q_2)d_2 x_2].$$

The incentive constraint for the governor is:

$$b_2 \geq (H_2 - L_2) - q_2 d_2 \delta_2 V_2,$$

where $\delta_2 V_2$ is the discounted value of the governor from keeping her position. Under the optimal contract the incentive constraint is binding. Setting the governor's utility of unemployment to zero, we obtain, in analogy to (1) in the basic model:

$$V_2 = p_2 b_2 + [1 - d_2(1 - p_2)(1 - q_2)]\delta_2 V_2,$$

from which it is possible to solve for $V_2(b_2, d_2)$ as in (2), and then solve explicitly for the optimal incentive scheme b_2 and d_2 selected by the king.

Thus, subject to additional parameter restrictions on x_2 and δ_2 that are analogous to those in the proposition, there exists a threshold $\hat{q}_2 > 1/2$ such that if village farming is sufficiently opaque to the king ($q_2 < \hat{q}_2$) the governor enjoys a carrot regime, in which she is autonomous in the sense that she is never dismissed, namely $d_{2c} = 0$. In this regime, the per-period revenue to the king is independent of the state of nature and is given by $\pi_{2c} = L_2$; the governor retains for herself the difference $b_{2c} = H_2 - L_2$ whenever the village state of nature is good, and zero otherwise.

On the other hand, when village farming is sufficiently transparent to the king ($q_2 > \hat{q}_2$), a stick and carrot regime prevails. Under this regime, the governor is dismissed whenever the king is led to expect high revenue based on his observed signal, but the governor reports that collected revenue is low. This occurs with probability $(1 - p_2)(1 - q_2)$. In this regime, following a similar derivation to the one in the basic model $d_{2s} = 1$ and $b_{2s} = (H_2 - L_2) - q_2 \delta_2 V_{2s}$, where:

$$V_{2s} = \frac{p(H_2 - L_2)}{1 - \delta_2(p + q_2 - 2pq_2)},$$

and the king's expected revenue is:

$$\pi_{2s} = (L_2 - m_2) + pq_2 \delta_2 V_{2s} - (1 - p)(1 - q_2)x_2.$$

The threshold transparency level $\hat{q}_2$ is determined by the implicit condition $\pi_{2s} = \pi_{2c}$. As in the basic model, the transparency threshold $\hat{q}_2$ increases with the cost of dismissal x_2 and decreases with the governor's discount factor δ_2.

As in the basic model above, the balance of power between the king and a provincial governor depends on the transparency of the provincial economy to the king. When local conditions are sufficiently opaque to the king, the intermediary governor enjoys substantial autonomy in that she pays a (relatively low) fixed tribute and retains her position forever. But if the transparency of the local provincial economy to the king is sufficiently high, then the governor is subject to dismissal and retains a relatively lower share of the revenue collected. If transparency is very high and the governor is very patient, she retains little beyond her minimal maintenance costs.