# Discipline and Selection: Explication of the Models in *Principled Agents*

November 3, 2008

## 1 Introduction

Besley's focus on *Principled Agents* is the tradeoff between discipline and selection in situations of political agency. In his models, some politicians want to serve the public and some politicians are self-serving. The latter will act like the former, and thus serve the public's interests, if it will help them stay in office and gather more rents. This is the electoral *discipline*. But the more discipline there is, the harder it is for voters to distinguish the public-minded politicians from those who merely seem public minded (i.e. to have effective *selection*). This matters to the voters because there is a limit to their ability to discipline politicians: at some point politicians will be free from discipline and the self-serving ones will make the voters pay.

In this note I carefully went through the models in the book and derived the expressions, most of which were presented as "straightforward" but some of which took some time to derive. My goal is to provide a more thorough account of the math and a more extensive discussion of the intuition behind the models. I didn't get a chance to cover all of the extensions to the models, so because I'm interested in models of transparency/endogenous political regulation I focused on those.

## 2 Besley's Chapter 3 Agency Model: Pooling Equilibria in a Model of Adverse Selection

The basic setup is that politicians have a type – sometimes called "good" or "bad," sometimes "congruent" or "dissonant" – that makes them value "doing what the voters want" vs. doing the opposite. Politicians are selected randomly from a pool of available politicians, the proportion $\pi$ of which are of the good type. After serving for one term, they face reelection; in this sense there is electoral accountability. The voters can either choose to reelect the incumbent or randomly draw a new politician from the pool. They make their choice based on an observed outcome, which depends on an unobserved action and an unobserved state. Still, voters in the basic model here know whether the action was consonant or dissonant based on their utility. But because politicians care about reelection they don't know if the consonant action was produced by a good politician who will continue to serve them well or a bad one who will screw them over in his lame-duck term.

The prevailing dynamic, similar to that in many principal-agent models, is that there are pooling equilibria, in which the dissonant politicians act like good ones in the first period, and separating equilibria, in which they act like bad politicians in the first term. Besley sets things up such that some dissonant politicians pool (ie act good) and some separate (ie act bad) in every equilibrium. The effect of pooling on voter welfare is the key interesting aspect of the model. When a bad politician acts like a good one, this is good for voters in the first period, but bad thereafter; because voters can't tell the difference, they will reelect the bad politician and get bad outcomes in the second period.

Besley sets up his model such that bad politicians face heterogeneous rewards for choosing the "bad" action and some politicians will always find it worthwhile to take the bad action and reveal their type. (In other words, there are no equilibria with complete pooling.) In this setup, it always weakly improves voters' utility to base their reelection decision on the consonance of the politician's action: if the voters choose to vote everyone out of office, or reelect everyone, all bad politicians will act badly whenever they are in office, so the probability of getting a good outcome is equal to the probability of randomly getting a good politician. By rewarding politicians for acting well in the first period, the voters can get good behavior in the first period from some bad politicians and, since some bad politicians will be weeded out, politicians in their lame-duck period will still be better than the average politician.

## 2.1 The basic agency model

### 2.1.1 Setup

There are two time periods, $t \in \{1, 2\}$. In each period, a politician makes a single decision, $e_t \in \{0, 1\}$. The payoff to voters and politicians depends on the state of the world $s_t \in \{0, 1\}$, which is only observed by the incumbent. If $e_t = s_t$, voters receive $\Delta$. (One could interpret this as saying that voters benefit if the politician takes the action that "fits" the situation.) Otherwise they receive nothing.

All politicians get benefit $E$ from serving in office (which might be thought of as ego rents or wage or both), but their policy-dependent payoff depends on their type. Congruent politicians (who make up a proportion $\pi$ of the pool of politicians) share voters' preferences, such that if they enact $e_t = s_t$ (as they always do) they get $E + \Delta$. Dissonant politicians get private benefit, or rent, $r \in [0, R]$ from picking $e_t \neq s_t$. (It is not explicitly stated here, but they do not get $\Delta$ from enacting the good policy.) This rent is drawn from a distribution with cdf $G(r)$ and known to the politician alone, before he chooses the policy. The mean of $r$ is $\mu$ and, in order to assure that the dissonant politician makes "bad" choices some of the time, Besley assumes that $R > \beta(\mu+E)$, ie, politicians with very high draws of $r$ will find it better to take the bad action in the first round rather than act good and get the rents in the next period. He introduces some very nice notation here, which he doesn't then use here very much, which says that the action of the politician in round $t$ is $e_t(s,i)$ with $s \in \{0,1\}$ and $i \in \{c,d\}$.

In the beginning of the game, a randomly-selected politician comes to power, $r$ is drawn, and the incumbent picks an action. Then voters observe the outcome and decide whether to re-elect the incumbent or draw another politician at random from the pool. If the incumbent is thrown out, we get a new politician for $t = 2$; if not, we get the same politician for a second term. In the most basic model the game ends here.

#### 2.1.2 Equilibrium

Since the game ends in period 2, politicians act in that period according to their type: $e_2(s,c) = s_2$ and $e_2(s,d) = 1 - s_2$. Consonant politicians always do $e_t = s_t$, so their behavior in the first period isn't interesting either. What is interesting is the action of the dissonant politician in period 1. The voters would reject any politician known to be dissonant, unless I suppose if $\pi = 0$, i.e. all politicians are dissonant, in which case they would be indifferent. So dissonant politicians in period 1 must consider whether to act in a consonant way now and hang on to act dissonantly in the next period, in which case they can expect to get $E + \beta(E + \mu)$ (because they will get paid in both periods but receive the rents (expected value: $\mu$) only in the second period), or to act in the dissonant way now and get $E + r_1$. The choice therefore depends on both one's first-period realization of $r$ and the mean of $r$. (If the politician were risk averse you would care about the distribution I guess. Or you could relabel $r$ the utility from deviating, in which case you are doing expected utility and all is fine.)

For the voters, the key issue is whether a politician observed to be acting in a consonant way is in fact a consonant politician. If voters receive $\Delta$ in the first period, they need to compare the probability that it was in fact a good guy vs just a bad guy acting good, because their next period outcome depends on that. Let's call $\lambda$ the probability that a bad politician will take the consonant action in the first period, ie $Pr(\Delta|d)$. (We'll derive that in a second.) From the voter's perspective the probability that a politician who delivers $\Delta$ is in fact consonant can be expressed as

$$\begin{aligned} Pr(c|\Delta) &= \frac{Pr(\Delta|c)Pr(c)}{Pr(\Delta|c)Pr(c) + Pr(\Delta|d)Pr(d)} \\ &= \frac{1 \cdot \pi}{1 \cdot \pi + \lambda(1-\pi)} \end{aligned}$$

This expression, which he calls $\Pi$, is weakly larger than $\pi$, which makes sense: if some dissonant politicians take the bad action in the first period (ie $\lambda > 0$) and some politicians are good (ie $\pi > 0$), then it has to be the case that you do better by sticking with a guy who gave you a good outcome in the first period rather than picking a politician at random, who will then implement his desired action with no accountability. ($\Pi = \pi$ if $\lambda = 1$, ie all dissonant politicians take the good action, which we have ruled out by assumption, or if $\pi \in \{0, 1\}$, ie all politicians are either good or bad, in which case it makes sense that you would be equally well off with a random politician as any politician who does any particular action, including delivering $\Delta$.)

Since we've shown that for normal circumstances voters will reelect any politician who delivers $\Delta$, we can now think about the dissonant politician's choice of whether to deliver $\Delta$ or not. The dissonant politician will act in a consonant way if $r_1 < \beta(E + \mu)$. Recalling that $G$ is the cdf of the rent distribution, the probability of this occurring is $\lambda = G(\beta(E + \mu))$, ie the proportion of bad politicians who receive a first period rent opportunity smaller than the expected gain from sticking around.

#### 2.1.3 Implications

Now we can think about the "quality of government," or how voters' utility depends on features of the model. We consider ex ante voter welfare in both periods.

$$\begin{aligned}
V_1(\lambda) &= \pi\Delta + (1-\pi)\lambda\Delta \\
&= \Delta(\pi + (1-\pi)\lambda) \\
V_2(\lambda) &= \pi\Delta + (1-\pi)(1-\lambda)\pi\Delta \\
&= \pi\Delta(1 + (1-\pi)(1-\lambda))
\end{aligned}$$

Voters unsurprisingly are better off in the first term if there are more good politicians around or more discipline among the bad ones. In the second period voter welfare is clearly decreasing in $\lambda$. The deriviative of the value wrt $\lambda$ in the first period is $\Delta(1-\pi)$, which is larger than that in the second period, which is $\beta\pi(1-\pi)\Delta$, which tells us that discipline is overall a good thing for voters in this model.[1]

### 2.2 Besley's extensions to the basic agency model

#### 2.2.1 Competitiveness of the political system

One extension that is interesting to me is the one in Besley 3.4.1 – "Polarization and competition." Here he applies a probabilistic voting model like the ones found in Persson and Tabellini, in which the voters have a preference for either the incumbent or the challenger, which changes the incumbent's action. The

[1]You can think of political moves that could be taken to change $\lambda$. Besley looks elsewhere at the wages politicians receive, $E$. You could also think about changing $G()$. Clearly with all of these things you'd get the same order of the $\frac{\partial V}{\partial \lambda}$ and then a $\frac{\partial \lambda}{\partial E}$ or whatever. So anything that increases $\lambda$, for example a mean shift downward in the distribution of $r$, would have this same effect of benefiting voters by increasing discipline more than it increases the corresponding decrease in selection.

basic idea here is that the issue we study is a "valence issue," ie an issue on which all the voters agree, and we can separate that issue from other issues on which parties are competing in the background. These other issues affect the voters though and affect the actions of the politicians on the valence issue.

In the model, $\omega$ fraction of the voters are assumed to be partisan voters who care about which party (A or B) is in office based on other issues (not part of the model). Besley operationalizes this by saying that partisan voters get utility $\phi > 0$ from having their preferred ideology in office. (For zero $\phi$ we have the model above; for very high $\phi$ the valence issue we study does not affect utility enough to sway partisan voters.) Things aren't interesting unless one party is favored by this partisan competition, which he operationalizes by saying that $\frac{1}{2} + \eta$ of the partisan voters favor party A.

There is noise in the votes of non-partisan voters: a popularity shock to the incumbent $\delta \sim U[-\frac{1}{2\xi}, \frac{1}{2\xi}]$ that affects all non-partisan voters the same way and an idiosyncratic shock $\iota \sim U[-\frac{1}{2}, \frac{1}{2}]$ that randomly affects each non-partisan voter.

Besley writes down the condition under which the incumbent beats a randomly selected challenger, calling it straightforward; I had to work this through for a while so I will write it here. We first try to determine the proportion of nonpartisan voters who will vote for party A, assuming the incumbent is from party A, as a function of the shocks in the model and the reputations of the incumbent vs a randomly selected politician (which, recall, are an equilibrium outcome of voter strategies and politician incentives). We start by noting that nonpartisan voter $i$ votes for a incumbent from party A who takes the good action when

$$\delta + \iota_i + \Delta(\Pi - \pi) > 0,$$

ie when the common shock and his own shock, added to the expected difference in utility in the upcoming period from retaining the incumbent vs picking a randomly selected politician, is greater than 0. If we write the cdf for $\iota$ as $F$ and assume as Besley does that it is symmetric (he actually assumes uniform), we know that the fraction of nonpartisan voters who vote for party A under this scenario will be $1 - F(-\delta - \Delta(\Pi - \pi)) = F(\delta + \Delta(\Pi - \pi))$, ie the proportion of nonpartisan voters with $\iota_i$ above $-\delta - \Delta(\Pi - \pi)$, which is the same as the proportion with $\iota_i$ lower than $\delta + \Delta(\Pi - \pi)$. What is that proportion? Besley assumes that $\iota$ is uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$, which means that $F(\iota) = \frac{1}{2} + \iota$. The proportion of nonpartisan voters who vote for a congruent incumbent from party $A$ can therefore be expressed as $\frac{1}{2} + \delta + \Delta(\Pi - \pi)$.

Now we can write out the conditions under which the congruent incumbent from party A beats a challenger. This happens when

$$\begin{aligned} \omega(\eta + \frac{1}{2}) + (1 - \omega)(\frac{1}{2} + \delta + \Delta(\Pi - \pi)) &> \frac{1}{2} \\ \omega\eta + (1 - \omega)(\delta + \Delta(\Pi - \pi)) &> 0, \end{aligned}$$

where the second statement comes from simplifying the first.

This is still written as a function of $\delta$, which is realization of a random variable. Rearrange the statement above as follows. $A$ wins when

$$\begin{aligned} \delta &> -\frac{\omega\eta}{1-\omega} - \Delta(\Pi - \pi) \\ &< \frac{\omega\eta}{1-\omega} + \Delta(\Pi - \pi). \end{aligned}$$

So viewed *ex ante*, the probability that the incumbent from party $A$ will win by taking the congruent action is $H(\theta + \Delta(\Pi - \pi))$, where $H(\delta)$ is the cdf of $\delta$ and $\theta \equiv \frac{\omega\eta}{1-\omega}$ can be thought of as a measure of partisan bias in favor of party $A$ (based on the proportion of partisans, $\omega$, and the extent to which they favor party $A$ , $\eta$). Given that $\delta$ is distributed uniformly on $[\frac{-1}{2\xi}, \frac{1}{2\xi}]$, the cdf is $H(\delta) = \frac{1}{2} + \xi\delta$. So the probability of $A$ winning under these conditions is written on page 126 as $\frac{1}{2} + \xi[\theta + \Delta[\Pi - \pi]]$, conditional on $\frac{-1}{2\xi} \leq \theta + \Delta[\Pi - \pi] \leq \frac{1}{2\xi}$ (i.e. as long as the system is not so uncompetitive, or the reputations so different, that the nonpartisan voters don't matter at all, in which case the probability of $A$ winning is 1 if the partisans favor $A$).[2]

Now, let's look at the condition on when a dissonant incumbent would choose the good action. We suppose first that the probability of winning lies between zero and one, ie we are in a system where the result depends on the realization of the popularity shock. I think we are also assuming here that the probability of winning if you do the dissonant action is not zero. The incumbent politician would decide by comparing the benefits as follows:

$$E + r_1 + H(\theta - \pi\Delta)\beta(\mu + E) \leq E + H(\theta + \Delta[\Pi - \pi])\beta(\mu + E),$$

where the key issue is whether the benefit this period is enough to offset the decrease in the probability of reelection. Note that the probability of winning by acting dissonantly is $H(\theta - \pi\Delta)$, which follows from the fact that the voters know that the incumbent is bad and expect to get $\pi\Delta$ next period if they choose a random politician.[3] This is not what Besley has (on page 126), but I'm fairly sure his is a mistake. He has the probability of winning by acting dissonantly as $H(\theta)$, which would imply that the voters expect the politician who has just acted dissonantly to be just as likely to deliver the consonant action in the next period as a randomly selected politician, ie they learned nothing about the type of the politician from observing him choose the dissonant action in the first

[2] Besley uses the notation $\sigma$ to refer to the probability of winning, but since this is always just a function of $\delta$ I prefer to write it in terms of the distribution of $\delta$.

[3] To see this, note that a non-partisan would vote for an incumbent from party A who takes the bad action if $\delta + \iota_i > \pi\Delta$, ie if the shock and personal preference outweighs the expected benefit from choosing a random challenger. The proportion of nonpartisan voters who would vote for a dissonant incumbent from party $A$ (assuming uniform idiosyncratic shocks) thus becomes $\frac{1}{2} + \delta - \pi\Delta$, which means that we can substitute $\delta - \pi\Delta$ for $\delta + \Delta(\Pi - \pi)$ in the above calculations. It also follows from thinking of $\Pi$ generally as a measure of reputation (a probability of being the good type, conditional on behavior), and here the reputation must be 0.

period. This has to be wrong. Assuming that the common shock is distributed uniformly, my expression can be restated as

$$\begin{aligned} r_1 &\leq \beta(\mu+E)[H(\theta+\Delta[\Pi-\pi]) - H(\theta-\pi\Delta)] \\ &= \beta(\mu+E)[\xi[\theta+\Delta[\Pi-\pi]] - \xi(\theta-\pi\Delta)] \\ &= \xi\Delta\Pi\beta(\mu+E), \end{aligned}$$

and the probability of a dissonant politician being discplined is thus $\lambda = G(\xi\Delta\Pi\beta(\mu+E))$. Quite reasonably, this says that electoral discipline goes up when the rewards of office are higher ($E$ and, for a given realized rent, $\mu$), when the politicians are more patient to receive those rewards ($\beta$), when "non-partisan" voters are more sensitive to the incumbent's action (where $\Delta$ represents the benefit to them from consonant action, $\xi$ indicates how much they care about this issue as opposed to other issues, and $\Pi$ indicates how strong of an indication consonant action is about future action). Note that neither $\theta$ nor $\pi$ appear in this expression (although recall that $\pi$ is a component of $\Pi$). Technically this is because $h$ is uniform and those factors enter into the probability of victory in the same way whether the incumbent takes the consonant or dissonant action. More intuitively, partisan advantage doesn't affect the degree of discipline under these assumptions because taking the consonant action improves the incumbent's probability of winning by the same amount regardless of the size of the partisan advantage.

With Besley, we go back and do away with the assumption that $\delta$ is uniformly distributed and suppose instead that $\delta$ has a density $h(\delta)$, which we require to be unimodal and symmetric about zero, to make the math easier. As before, an incumbent from party $A$ who takes the congruent action wins with probability $H(\theta+\Delta(\Pi-\pi))$ (where we assume again that the policy and competitiveness are not such that we are at a corner where the probability is zero or one). As before, a dissonant politician will take the congruent action if

$$r_1 \leq \beta(\mu+E)[H(\theta+\Delta[\Pi-\pi]) - H(\theta-\pi\Delta)].$$

Since we no longer assume that $h$ is uniform, we cannot simplify using the uniform CDF. The probability of the above being true, $\lambda$, can be written as

$$\lambda = G\Big(\big(H(\theta+\Delta[\Pi-\pi]) - H(\theta-\pi\Delta)\big)\beta(\mu+E)\Big).$$

In words, the probability of a dissonant taking the good action is equal to the probability of the realized rent being smaller than the expected benefit from acting congruently. Now we ask how this expression varies with $\theta$, ie how the extent of discipline varies with pro-A bias. We know $G$ is monotonically increasing, because it's a cdf. So we ask how its argument changes with $\theta$:

$$\begin{aligned} \frac{\partial\big(\beta(\mu+E)(H(\theta+\Delta[\Pi-\pi]) - H(\theta-\pi\Delta)\big)}{\partial\theta} &\propto \frac{\partial(H(\theta+\Delta[\Pi-\pi])}{\partial\theta} - \frac{\partial H(\theta-\pi\Delta)}{\partial\theta} \\ &= h(\theta+\Delta[\Pi-\pi]) - h(\theta-\pi\Delta). \end{aligned}$$

This is the same as Besley's expression once we fix the error noted above.[4] If, as was assumed in the earlier example, $h$ is a uniform density, this is obviously zero and the incumbent's choice of whether or not to take the congruent action does not depend on competitiveness. In the case where $h$ is symmetric and unimodal and the partisan advantage $\theta$ is large enough relative to $\Delta$ and $\pi$,[5] the above expression is negative, indicating that discipline goes down when partisan bias goes up. With a unimodal popularity shock the benefit to the incumbent from discipline depends on the extent of the partisan advantage; intuitively, this is because the probability that taking the consonant action will affect the election outcome becomes lower as the safeness of the seat increases. In equilibrium therefore there will be less discipline when the system is less competitive.

Besley's proposition 2 is thus that, when the incumbent has an advantage, decreasing the competitiveness will also decrease his discipline. As shown above, lower discipline means lower voter utility in the first period but higher in the second; the first-period effect outweighs the second such that the net effect of lower competition is to decrease welfare for the non-partisan voters. (Since we are saying that the partisan voters have policy preferences it's a little harder to generalize about the net effect on welfare for all voters.)

### 2.2.2 Information and accountability (3.4.2)

Besley explores the role of information by limiting one kind of information voters can receive and creating another new kind of information. First, the limit: voters do not always know whether the politician has taken the good action or not. In particular, with probability $1 - \chi$ the voters don't know whether they received $\Delta$ or not.[6] On the other hand, voters in this version of the model have additional insight into the politician's type. In particular, voters directly observe the politician's type before the reelection with probability $\tau$.

In order to progress, he assumes that if no information is revealed (ie the voters don't find out the type and don't know the state of $\Delta$) the voters will reelect the incumbent. I suppose this will be justified later. Note though that when no information is revealed this mean several things: a good politician taking the good action, a bad politician taking the good action, or a bad politician taking the bad action.

Actions in the second period are not changed, nor would a congruent politician act dissonantly in the first period. What might change, as is always the case with this model, is the behavior of the dissonant politician in the first period. If the dissonant politician takes the dissonant action, the voters come to know that he is dissonant either by finding out his type or the action, which happens with probability $\tau + (1 - \tau)\chi$. Since we assume that voters elect unless

[4] Note that this is not exactly right either, though, since $\Pi$ depends on $\lambda$ and therefore $\theta$.

[5] The specific condition is that $\theta > \frac{\Delta(2\pi - \Pi)}{2}$. Besley's condition, based on the erroneous expression for the probability of winning by taking a dissonant action, is $\theta > 0$.

[6] At first I thought this was a very awkward way to model transparency, since if they don't receive $\Delta$ they shouldn't care about the act; but in fact they still feel the effects in the second period, so while this complicates the welfare analysis a little it's not too bad.

they know the politician is dissonant, we know that a dissonant politician acting dissonantly will be reelected with probability $1-(\tau+(1-\tau)\chi)=(1-\tau)(1-\chi)$. (Previously he would be reelected with probability zero.) If he takes the congruent action, he will be reelected with probability $1-\tau$, ie unless voters find out about his type directly. So we have that a dissonant politician will take the consonant action if

$$\begin{aligned} E+r_1+(1-\tau)(1-\chi)\beta(\mu+E) &\leq E+(1-\tau)\beta(\mu+E) \\ r_1 &\leq (1-\tau)\chi\beta(\mu+E) \end{aligned}$$

which implies a probability of disciplined action $\lambda=G((1-\tau)\chi\beta(\mu+E))$. This expression suggests that discipline depends on the two types of information differently. The greater the public's understanding of the actions politicians take $\chi$ the more likely dissonant politicians are to be disciplined by electoral pressure. On the other hand, the greater the powers of the public to discern politicians' true types $\tau$ the less discipline there will be. Previously we saw that more discipline improves voter welfare (because the improvement in first-period action compared to sincere behavior outweighs the decrease in second period action relative to the randomly selected politician). Here we have the added wrinkle that these information factors enter voter utility not just through the effect on discipline (ie on the policy choices of dissonant politicians) but also directly, since when the voters observe the type or action they can do something about it. So we'll need to look more carefully at how information affects voter welfare here.

Voter welfare can be written as

$$\begin{aligned} W(\lambda;\tau;\chi) &= \pi(\Delta+\beta\Delta)+(1-\pi)(\lambda(\Delta+\beta(\tau\pi\Delta+(1-\tau)0))+ \\ &\quad (1-\lambda)\beta(\tau\pi\Delta+(1-\tau)\chi\pi\Delta+(1-\tau)(1-\chi)0)) \\ &= \pi(\Delta+\beta\Delta)+(1-\pi)(\lambda(\Delta+\beta\tau\pi\Delta)+(1-\lambda)\beta(\tau\pi\Delta+(1-\tau)\chi\pi\Delta)) \end{aligned}$$

In words, with probability $\pi$ we have a consonant type; we will never get information that will cause us to reject him, so we get $\Delta+\beta\Delta$ out of him. With probability $1-\pi$ we get the dissonant type. With probability $\lambda$ he takes the consonant action, so we get the benefit in this period and no future benefit unless we observe the type directly in this period (in which case we get the future benefit with probability $\pi$); with probability $1-\lambda$ he takes the dissonant action, so we get no benefit this period and in the next period get a benefit if we observe the type and get a good type next time (with probability $\tau\pi$) or we don't observe the type but we do observe the action/benefit and we get a good type next time (with probability $(1-\tau)\chi\pi$). The expression can be simplified as

$$W(\lambda;\tau;\chi) = \Delta(\pi+(1-\pi)\lambda+\beta\pi(1+(1-\pi)(\tau+(1-\lambda)(1-\tau)\chi))),$$

which Besley writes down as

$$W(\lambda;\tau;\chi) \quad = \quad \Delta(\pi + (1-\pi)\lambda + \beta\pi(1 + (1-\pi)(1-(1-\tau)(1-\chi(1-\lambda))))),$$

which is somewhat more cumbersome but makes it a little easier to calculate partial derivatives. Now we can look at the partial derivatives of voter value with respect to these information parameters. We have:

$$\begin{aligned}\frac{\partial W}{\partial \chi} &= (1-\pi)((1-\beta\pi(1-\tau)\chi)\frac{\partial \lambda}{\partial \chi} + \beta\pi(1-\tau)(1-\lambda))\Delta \\ \frac{\partial W}{\partial \tau} &= (1-\pi)((1-\beta\pi(1-\tau)\chi)\frac{\partial \lambda}{\partial \tau} + \beta\pi(1-\chi(1-\lambda)))\Delta,\end{aligned}$$

which are obviously quite similar to each other. (Besley does not write out $\frac{\partial W}{\partial \chi}$, but he does write out $\frac{\partial W}{\partial \tau}$ and has the same thing I do, although he appears to have made a typo and written $\lambda$ instead of $\Delta$ as the last term in the expression.) The first expression is positive, since $\frac{\partial \lambda}{\partial \chi}$ is positive, indicating that knowing about the action increases welfare unambiguously. (There would certainly be diminishing returns here, even if G were linear (eg if $g()$ was uniform), because at higher levels of information $\frac{\partial W}{\partial \chi}$ gets smaller.) The story with $\frac{\partial W}{\partial \tau}$ is ambiguous, indicating that knowing more about the true type of the politician does not necessarily increase voter welfare. The first term will be negative and the second positive, so the whole expression will be negative for small $\pi$, indicating a corrupt pool of politicians, small $\beta$ (indicating that voters and politicians are impatient), or high $\lambda$ (indicating relatively high discipline). Besley explains the result on $\pi$ by saying greater non-policy information about incumbents "reduces welfare when the fraction of dissonant politicians is sufficiently large . . . because the selection benefits of better information are negligible in this case – it is highly likely that one dissonant incumbent will be replaced by another – and yet there is a negative disipline effect." In other words, as information about types gets better, voters get a better chance to throw bad incumbents out of office, but the benefit from doing so is clearly smaller when the pool of politicians (and thus the expected quality of the replacement) is worse; at the same time, knowing more about types discourages bad incumbents from taking the good action because they know they may not be rewarded with another term even if they do take the good action. So as we consider worse and worse politician pools, at some point the selection benefit gets small enough that it no longer outweighs the decrease in discpline from knowing more about types, and you would actually be better off knowing less about politician types because this would again encourage bad politicians to pretend to be good in order to get reelected. This is quite an interesting, counter-intuitive result until you've thought about this discipline/selection tradeoff for a while.

### 2.2.3 Reality check

It's worth thinking about whether this agency model captures important features of real political life. Here are key features of the model and some reflection on their versimilitude:

- **New politicians are chosen at random from a pool**. This does not sound realistic but we can think of it as abstracting from the idea that new politicians are of basically the same quality, and less is known about them than about incumbents.
- **There are no positions taken here, so everything is about politician type.** It seems plausible that politicians differ in their taste or capability for taking actions that benefit themselves at the expense of the public.
- **In the second term, the politician has no constraints**. Besley and Case's work does suggest that politicians act differently in their second terms (more taxes and spending) in a way that suggests that discipline is lower in the second period. More generally we could think of the model as a situation in which voters choose based on limited information whether to give the politician powers in an area they cannot control. Saying that there is a lame-duck period is one way to capture the fact that there are limits to electoral accountability. Here the lack of accountability in the second period comes from the fact that time (or political careers) ends, and voters act based on what politicians deliver (not what they say they will deliver). Perhaps the setup could be reconceived to mean that there is a part of the political process that is scrutinized, and part that is not; ie part that is liable to discipline and part that is not. The discipline dynamic would thus be that politicians may act better in the scrutinized part in order to get the benefits of the unscrutinized part. Transparency in that model might amount to decreasing the size of the unscrutinized part, from the perspective of the politician. When we take it out of an electoral setting, though, it's hard to motivate the fact that voters have the discretion to deny politicians access to the scrutinized part, if not through reelection.

# 3 Besley's Chapter 4 Model: Adverse Selection and Moral Hazard

Previously we had hidden types for the politicians, but the voters knew perfectly if the politician had taken the good or bad action (except in the "transparency" extension). Besley's chapter 4 explicitly combines adverse selection and moral hazard by setting things up such that the voters cannot deduce the type from the action even when the dissonant politicians are getting rents. The dynamic of the tradeoff between discipline and selection remains in the foreground, but some of the comparative statics change based on the ability bad politicians now have to get rents and not be thrown out.

## 3.1 Setup

The government decides three things in each period: the level of "productive" public spending $G_t$, the level of public spending devoted to private ends $s_t$, and the level of taxation $x_t$, which is constrained to be between 0 and $X$. This is how he presents it, although as it turns out it is more useful to think of $G_t$ as the amount of public goods (ie not spending). In each period the unit cost of providing public goods is $\theta \in \{L, H\}$, with $H > L$ and $Pr(\theta = H) = q$. This unit cost is observed by the politician but not by the voters. The budget must balance, so $x_t = \theta G_t + s_t$. This sets up a situation where, when $x_t > LG_t$, the voters don't know if this is because $\theta = H$ or because the politician is stealing $s_t = (H - L)G_t$.

Here again we have politician types, $i \in \{b, g\}$. A good politician does not steal and instead picks the optimal amount of public goods. The bad politician behaves strategically, trying to maximize his expected discounted stream of rents $s_1 + \beta\sigma s_2$, where $\sigma$ is the probability of being reelected. Again, the model lasts two periods, so there is no interesting behavior in the second period, and the key issue is how the bad types behave in the first period.

Compared to the earlier model, the key feature is that here the voters don't know whether the politician took the "correct" action or not. Before, the voters didn't know the type of the politician and no one knew the specific rents that would be available from dissonant action. But when the politician acted dissonantly the voters knew, and only bad politicians would act dissonantly. So the issue was just whether bad politicians would act good sometimes in order to be reelected. Here we still have the opportunity for bad politicains to pool with good politicians, but we add the wrinkle that, because voters don't know the state of the world (cost of public goods), politicians might be able to get rents in equilibrium without the voters knowing they're bad.

## 3.2 Three scenarios

### 3.2.1 Pure adverse selection

As a starting point, we assume that politicians do not behave strategically at all. Good politicians take the good action and bad ones steal a pre-specified amount. $G$ is assumed to be fixed at some commonly known level $\bar{G}$. Bad politicians steal an amount that manages to make their type ambiguous in the case where $\theta = H$; this amount is $(H-L)\bar{G}$, since $L\bar{G}+(H-L)\bar{G} = H\bar{G}$. Given these assumptions, voters will observe only three levels of taxation, $x \in \{L\bar{G}, H\bar{G}, (2H-L)\bar{G}\}$. The first one only happens when $\theta = L$ and the politician is good (and therefore not stealing anything), and the last happens when $\theta = H$ and the bad politician is stealing the specified amount. From the perspective of the voters, then, the outcome in which they are confused about the politician's type is $H\bar{G}$.

By Bayes' Rule, The probability that the politician is good, conditional on seeing $H\bar{G}$, is

$$Pr(g|H\bar{G}) = \frac{Pr(H\bar{G}, g)}{Pr(H\bar{G}, g) + Pr(H\bar{G}, b)} = \frac{q\pi}{q\pi + (1-q)(1-\pi)},$$

ie the probability that the politician is good conditional on seeing $H\bar{G}$ is the probability of seeing $H\bar{G}$ and $g$ divided by the probability of seeing $H\bar{G}$, under the different possibilities of politician quality. We see $H\bar{G}$ when the politician is bad if the true strate is $L$, which happens with probability $1-q$, and the politician is bad, which happens with probability $1-\pi$. I confirmed that the expression above is greater than $\pi$, the probability of a randomly selected politician being good, if $q > \frac{1}{2}$. At first it is surprising that this does not depend on $\pi$: we should reelect the politician upon seeing $H\bar{G}$ if the probability of $H$ is above $\frac{1}{2}$. I guess intuitively this is because, while the probability that a politician who gives you $H\bar{G}$ is good certainly depends on $\pi$, so does the probability that his replacement will be good, so apparently they cancel each other out. Because we are comparing with the probability that a random guy is good, we end up asking not "How likely is this to be a good guy?" but rather "Well, how likely is it that we are in a state that means this a good guy vs a state that means this is a bad guy?"

#### 3.2.2 Pure moral hazard, ie a pure screening model

Here we get a screening model that is straight out of MWG. There are now no intrinsically good politicians, so everyone wants to get as much rent as possible. (In the labor-market screening examples this is like the case where $G$ is effort/tasks/education and $s$ is wage, and $\theta$ is the hidden productivity of the workers. In those cases the workers are not assumed to be "bad" but rather quite naturally they care about wages and want to exert little effort.) The voters in this model commit to a re-election rule $\sigma(G, x)$ based on the performance of the politician. The politicians therefore want to choose levels of stealing to maximize $s_1 + \beta\sigma(G, \theta G + s_1)s_2$.

Voters need to let politicians get enough rent in the first period to make it worth holding back instead of taking all $X$ available. This threshold level of rent satisfies $\hat{s}(\sigma) + \sigma\beta X = X$, where it is written as $\hat{s}(\sigma)$ because the amount they need depends on their likelihood of being reelected. Since the politician doesn't care about $\theta$ per se, the voters have to give him this amount regardless of the state of the world. And in order to minimize this necessary amount of stealing $\hat{s}$ they should re-elect the politician with certainty.

Then we can consider what levels of public goods and taxes the voters will require. As in the classic screening problem, the voters' best strategy is to require a combination of taxes and spending that satisfy the ICC such that the politician does not pretend that the state is $L$ when it is actually $H$ and the PC that keeps the politician from stealing everything in the first period when the state is $H$. Figure 4.1 depicts this situation graphically. The flatter line is the PC: $s_H = x_H - HG_H = (1-\beta)X$. If the voters know that the state is $H$

they will do best by choosing the point of tangency between the PC and their indifference curves. But otherwise they can do better by accepting lower taxes and spending in the $H$ state in order to get some more utility in the $L$ state. If they go for the more austere government in the $H$ state, they can get more public goods in the $L$ state; they can thus make the $L$ situation more generous for the politician (and better for themselves) while discouraging the politician from dissembling and stealing all the surplus while pretending that the state is $H$. This is the standard screening model, as explained in MWG.

#### 3.2.3 Combining the two

Here, as in the Chapter 3 model, we have types of politicians and they can choose their actions. In the second period, $s_2 = X$ for $i = b$ and $s_2 = 0$ for $i = g$.

The voter's reelection decision depends on comparing the posterior probability that the politician is good to the prior probability that a random politician is good. The good politician will only spend $x \in \{x_L^*, x_H^*\}$, ie the optimal, no-stealing amounts for the low and high levels of public goods cost. Therefore a politician who chooses any other amount is known to be bad. The bad politician in equilibrium can therefore choose only $x \in \{x_L^*, x_H^*, X\}$, ie the same things the good politician would choose plus the maximum stealing amount. Let's consider each of these. If the bad politician chose $x_L^*$, that would imply that $\theta = L$ and $s_1 = 0$. But the bad politician would never choose $s_1 = 0$, because the most he could get by exerting that restraint would be $\beta X$ and he could do better by just choosing $X$ and being voted out of office. Thus $Pr(g|x_L^*) = 1$. We also know that $Pr(g|X) = 0$. (This is like the Chapter 3 model, where $Pr(g|0) = 0$.) The more complicated situation is when $\theta = L$, and the bad politicians can pretend that $\theta = H$ and pocket $(H - L)G_H^*$. As before, we write the probability of discipline in this case

$$\lambda = Pr(x = x_H^* | \theta = L, i = b).$$

The posterior probability that a politician who delivers $x_H^*$ is of the good type is then

$$\begin{aligned} Pr(g|x_H^*) &= \frac{Pr(x_H^*|g)Pr(g)}{Pr(x_H^*|g)Pr(g) + Pr(x_H^*|b)Pr(b)} \\ &= \frac{Pr(H)Pr(g)}{Pr(H)Pr(g) + \lambda Pr(L)Pr(b)} \\ &= \frac{\pi q}{\pi q + \lambda(1-\pi)(1-q)}. \end{aligned}$$

The voters should re-elect a politician who produces $x_H^*$ if $Pr(g|x_H^*) > \pi$, which is true when $\lambda < \frac{q}{1-q}$. Intuitively, we reelect if the discipline induced in equilibrium is low enough, or the probability that we actually are in the $H$ state high enough, that we think this guy is more likely to be a good politician than a randomly selected politician would be. (Compare this to the pure adverse selection model above, where politicians do not behave strategically; there it depended only on $q$.)

There are three equilibria.

- One is the **hybrid equilibrium,** where $\lambda = \frac{q}{1-q}$, such that the voters are indifferent about reelecting a politician who delivers $x_H^*$, and voters reelect with probability $\sigma = (X - \hat{s}(\mu))/(\beta X)$. ($\hat{s}(\mu)$ is defined as the stealing amount that makes it ambiguous whether we're in $H$ with a good politician or $L$ with a bad one.) This happens when the low state is more likely than the high state and the amount you can steal $\hat{s}(\mu)$ is large enough to make you want to wait (ie $\hat{s}(\mu) \geq (1-\beta)\mu$). Intuitively in this case the probability of reelection for delivering the ambiguous outcome is such that bad politicians are indifferent between doing that and stealing everything so they do some of each. *(I could be a little clearer here.)*

- The **pooling equilibrium** where bad politicians act good and are reelected. This exists if and only if $q \geq \frac{1}{2}$ (because that is required for $\lambda$ to be 1 and smaller than $\frac{q}{1-q}$) and $\hat{s}(\mu) \geq (1-\beta)X$. Intuitively here we have discipline because it's worth it for the bad politicians, considering that they can trick the voters because the voters believe a politician acting this way is more likely to be good than a randomly selected politician.

- The **separating equilibrium** where bad politicians act bad and are not reelected. This exists if and only if $\hat{s}(\mu) \leq (1-\beta)X$, ie if a bad politician doesn't find it worth pretending to be good in order to get the benefit in the next period.

## 3.3 Extensions

### 3.3.1 Information provision

Besley fairly briefly considers the role of information provision in the chapter 4 model. Here again the general point is that having more information about what the politicians are doing will improve selection but discourage discipline. He models this by providing a way in which the voters can learn about the true cost of public services $\theta$. He supposes that this becomes possible before the election with probability $\xi$. Now the politician has a different calculus of whether it makes sense to be disciplined or not. Specifically, he now will be disciplined if $s(\mu) + (1-\xi)\beta X \geq X$, ie. if $\hat{s}(\mu) > (1-(1-\xi)\beta)X$, which means that he needs to be able to steal more than previously. This would mean that for some parameter values we would move from a pooling to a separating equilibrium.

What is the welfare effect on the voters of such a shift toward the separating equilibrium? Using expression 4.2 on page 189, Besley is able to express the welfare of voters in terms of a baseline welfare when no one is reelected, a discipline effect $\Delta(\mu)$ that comes from bad politicians acting good, and a selection effect $\Sigma(\mu)$ that comes from having good incumbents instead of bad ones. The latter terms are preceded by coefficients that indicate how likely we are to be in a situation where that benefit takes place, i.e. how likely we are to have a bad politician and be in the low-cost state and have the politician exert discipline, and how much more likely we are to have a good politician in the second period than we would if we picked at random. In the state where $\theta = L$, we get selection in the separating equilibrium (ie no bad politicians make it to the second period) and discipline in the pooling equilibrium (ie no one acts bad in the first period). (We get selection when the state is $H$ in both cases.) Accordingly, the difference in welfare between the separating and pooling equilrium reflects the probability of being in the $(b, L)$ state and the difference in welfare that results. In moving from pooling to separating you get the welfare benefit

$$(1-\pi)(1-q)(\beta\pi\Sigma(\mu) - \Delta(\mu)),$$

where the inside term reflects the fact that the benefits of selection are discounted and depend on getting a good politician, while the discipline happens now. The key point here is that the expression is positive or negative (ie increasing the probability of discovering the bad politician and thus discouraging pooling is a good or bad thing) depending on the size of the selection versus the discipline effect on welfare (and $\beta$ and $\pi$). This in turn depends on the utility voters get in the different states of the world and the probability of those states happening.

The term above is basically the derivative of the voters' welfare expression with respect to $\lambda$ (see this on page 190), although the sign is reversed. In other words, increasing discipline increases welfare if the discipline effect outweighs the selection effect, ie if $\Delta(\mu) - \beta\pi\Sigma(\mu)$. Increasing scrutiny of the true state decreases discipline.

The point is thus the same as that in Chapter 3 and in the other extensions here: a fundamental tradeoff between selection and discipline. Convincing bad politicians to act like good politicians improves outcomes now but makes voters worse off in the lame-duck future, when politicians face no discipline. Put another way, identifying bad politicians to avoid pain later is only possible if we induce the bad politicians to take actions that hurt voters.

A key question to address here is, "Why is discipline *per se* always good in Chapter 3 but not in Chapter 4?" I believe the reason is that, due to the changing states of the world and their effect on voter welfare, the selection and discipline effects are of different sizes in this chapter, whereas they were the same in the earlier chapter. In the first model, the state did not affect the welfare benefit of congruent action; all congruent action produced utility $\Delta$. As a result, the expected welfare benefit from convincing a bad type to act good (the benefit of discipline) was $\Delta$, as was the expected welfare benefit from getting a good type rather than a bad type in the second period (the benefit of selection). In the second model, the discipline benefit is the difference in the voter welfare in the low-cost state when the politician makes the welfare-maximizing choice vs when the politician steals everything, while the selection benefit is the *expected* voter welfare under a good politician vs the welfare when the politician steals everything. In the first model, these basic benefits are of the same magnitude but the latter was of smaller order ($\beta\pi(1-\pi)$, because it's discounted and we get the benefit if we get a bad politician in the first period and a good one in the second, compared to $1-\pi$ because we get the benefit if we get a bad politician in the first period). In the second model, the relative size of the two effects depends on a number of parameters, so we can't say anything general about the sign of the partial derivative.

As a confirmation of the importance of the interaction between state and politician type in determining benefits, note that if we have the state always be $L$ (ie $q = 0$), such that the game simplifies to the one where bad politicians either steal everything or steal nothing in order to stay in office, the situation is back to that of the first model, if we express as $\Delta$ the difference between the utility we get with the good politician's choice and when the bad politician steals everything.