

Discussion Papers
Department of Economics
University of Copenhagen

No. 15-11

How Jeremy Bentham would defend against coordinated attacks

Ole Jann

Christoph Schottmüller

Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K., Denmark

Tel.: +45 35 32 30 01 – Fax: +45 35 32 30 00

<http://www.econ.ku.dk>

ISSN: 1601-2461 (E)

How Jeremy Bentham would defend against coordinated attacks*

Ole Jann

University of Copenhagen

Christoph Schottmüller

University of Copenhagen and TILEC

August 12, 2015

Abstract

How can a single player defend against the threat of a coordinated attack by a group? For example, how can a central bank defend a currency peg against speculators, a government against a revolution or a prison warden against a breakout? Bentham (1787) proposed an innovative prison concept based on information asymmetries – the “panopticon” – as an answer to this question. We consider different information structures in a stylized model of a prison, in which a warden chooses a costly guard level with the goal of avoiding breakouts. Successful breakouts require coordination among prisoners. We show that the information structure corresponding to the panopticon often performs best, especially if there are many prisoners.

Keywords: panopticon, coordination games, global games, transparency

JEL classification: D23 (Organizational Behavior), D74 (Conflict, Revolutions), D82 (Asymmetric and Private Information), E58 (Central Banks and Their Policies), Z13 (Economic Sociology), F31 (Foreign Exchange)

*Address to both authors: Department of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, Denmark; email: ole.jann@econ.ku.dk and christoph.schottmueller@econ.ku.dk. We are grateful for helpful comments from seminar audiences at Bocconi University, the University of Copenhagen and Tilec.

Morals reformed – health preserved – industry invigorated – instruction dif-fused – public burthens lightened – Economy seated, as it were, upon a rock – the gordian knot of the Poor-Laws not cut, but untied – all by a simple idea in Architecture! (*Bentham, 1787*)

1. Introduction

We analyze situations in which a single player is in conflict with a group of others, and the group members’ actions are strategic complements. Consider, for example, a government threatened by a revolution: Each potential revolutionary has to decide whether to show up for a demonstration, and larger demonstrations are more likely to succeed – but no one wants to be the only one to show up. A speculative attack on a currency peg requires the participation of many speculators – but if the attack fails because not enough speculators participate, those who participated will lose money.

In each of these cases, the single player would like to prevail with a minimum use of resources (security forces, currency reserves) by discouraging the group from acting in the first place. In this paper, we consider how he can accomplish this goal by choosing the right information structure. The information structure determines which information about his own strength will be revealed when he chooses a costly strength level at a later stage. Our surprising result is that in many situations, complete secrecy is optimal. That is, the single player foregoes the option to publicly commit himself to a strength level. Complete secrecy mirrors the idea of the “Panopticon” proposed by Bentham (1787) – an innovative prison concept in which prisoners were to be kept unable to see the guards as well as separated from each other.

The general problem that we consider has many applications, some of which we discuss later in the paper. Our main analysis concentrates on a succinct and graphic example close to Bentham’s idea: The question of how to construct a prison.¹ The prison warden faces a trade-off, as guards are costly but more guards offer more protection. The prison design allows a choice over how much information about the guard strength is available to the prisoners. This can make coordination among individual prisoners, in the absence of institutions that allow for explicit coordination, easier or harder. Ideally, the prison warden would prefer to maintain order in the prison and prevent revolts and breakouts while using a minimum of guards. The optimal prison design will exploit the prisoners’ coordination problem in order

¹Bentham tried to construct the actual Panopticon after his plans, using considerable time on the purpose while trying to convince successive governments of the idea. Unlike him, we mostly see the prison as a metaphor for the mechanisms we want to analyze. Taking our formal model as a practical guide to prison construction is done at the reader’s own risk.

to prevent them from revolting.

Bentham proposed that the isolation of the prisoners, together with their lack of knowledge of how many guards (if any) were on duty, would make coordination and thus a successful revolt impossible.² Through the lense of game theory, this argument appears unconvincing. Rational prisoners should be able to implicitly coordinate, and in equilibrium they should be able to infer the choice of the prison warden about guard strength. We find, however, that Bentham’s intuition plays out: In a large prison, where prisoners have no information about guard strength before independently choosing whether to revolt or not, there is only one equilibrium in which the warden randomizes between minimal guard levels and prisoners almost never revolt.

We compare four different information structures, also shown in table 1: (1a) Prisoners can observe the guard level and coordinate (“benchmark model”). (1b) Prisoners cannot observe the guard level but can coordinate (“benchmark model”). (2) Prisoners can observe the guard level but face a coordination problem (“infection model”). (3) Prisoners cannot observe the guard level and face a coordination problem (“panopticon”).

In cases (1a) and (1b), preventing a revolt is only possible when choosing the guard level such that a revolt by all prisoners would not be successful. In (2), “a union of hands” is required for a successful revolt for any intermediate guard level. As the actions of the prisoners are strategic complements, there are two equilibria in the prisoners’ subgame (after the warden has chosen an intermediate guard level): All prisoners revolt, or none does. One of these subgame equilibria (the successful revolt) is preferred by the prisoners, but in this equilibrium each of them puts himself at the mercy of the others – he does not want to be caught as the only one revolting. Following the global games literature, we assume that the prisoners, being isolated from each other, do not achieve *common* knowledge of the guard level. Without common knowledge, their higher-order beliefs will then be infected (Rubinstein, 1989; Carlsson and van Damme, 1993): ‘I believe that a revolt can be successful, but what if the others think that I think that it cannot? Then they would not revolt, and neither should I.’ This infection makes it possible to reliably prevent a revolt with a much lower guard level than in the benchmark model. While the number of guards needed to deter revolts still rises linearly in the number of prisoners, the factor is usually much lower than 1.

Finally, in the fourth model, the panopticon, it is not immediately obvious what kind of equilibria there are. Knowing that the guard level will not be observed, the warden has an incentive to choose a low guard level, but that will make him very vulnerable to revolts by

²Bentham (p. 46): “Overpowering the guard requires an union of hands, and a concert among minds. But what union, or what concert, can there be among persons, no one of whom will have set eyes on any other from the first moment of his entrance? ... But who would think of beginning a work of hours and days, without any tolerable prospect of making so much as the first motion towards it unobserved?”

		Guard level observable	
		Yes	No
Coordination problem between prisoners	No	(1a) Benchmark	(1b) Benchmark
	Yes	(2) Infection	(3) Panopticon

Table 1: The four information structures we consider.

even a few prisoners. Especially if there are many prisoners, it might seem sensible to always set a sufficiently high guard level to prevent substantial revolts.

Instead, we find that – if the number of prisoners is large – there is a unique equilibrium in mixed strategies in which the warden randomizes between the lowest possible guard levels, and each prisoner randomly chooses whether to revolt or not. The individual probability of revolting and the probability of a successful breakout are very small if the number of prisoners is large. This guarantees that revolts can be prevented almost surely with just one guard, as Bentham predicted. No other equilibria exist – neither pure nor mixed, symmetric or asymmetric. There can be no pure-strategy equilibria since breakout would then occur in equilibrium with probability 0 or 1. In the former case, either the warden would want to deviate to a lower guard level or the prisoners would want to deviate to attempting a breakout; in the latter case, the warden would want to deviate to a higher guard level.

In mixed equilibria, the distribution of the number of revolting prisoners matters for the chance of a successful breakout, which is the determining factor in both the warden’s and the prisoners’ strategic considerations. Technically, our argument is closely related to the law of large numbers and the tail bounds of probability distributions. If there are many prisoners, all of whom revolt with some probability, the actual number of revolting prisoners is more and more closely distributed around the expected number of revolting prisoners. If this expected value is close to the actual guard levels so that a successful revolt becomes likely, the warden will want to increase the guard level. If the guard level is much higher than the expected value, the probability of successful revolts is too low to induce prisoners to revolt. Together, these effects preclude the existence of any equilibria in which the warden chooses a guard level higher than the two minimal levels with positive probability. This fact makes the panopticon the optimal information structure for large groups of prisoners, where it performs far better than the other structures.

This result is close to what Bentham proposed. He envisioned the impossibility of a ”concert among minds” to such a degree that prisoners would not even think about revolting together with other prisoners, and would simply concentrate their thinking on the possibility of being watched and disciplined. If the number of prisoners is large, our model exhibits the same property: For any prisoner, the probability that any of the other prisoners will revolt

is close to zero, and the prisoner de facto finds himself in a game only between himself and the warden – where the warden chooses a mixing between having one guard and having no guards at all that just assures the prisoner’s docility. This result also has significance in understanding the social science literature that followed.

In the 230 years since Bentham, many scholars have interpreted the panopticon as a metaphor for modern society. Most prominently, Foucault (1975) points out that panopticism, a system in which individuals self-discipline because of the omnipresent possibility of being disciplined, has made modern society possible. Order is no longer maintained by overwhelming force or a contest of violence between those opposing and those defending it, as in our first model. Instead, the docility of individuals allows for cost-saving minimal enforcement.³ This allows for the establishment of organizations, firms, schools in which individuals have internalized the rules and behave in the desired way without constant supervision. It was this “accumulation of men” (p. 220) that, besides the accumulation of capital, allowed the industrial take-off of the early 18th century. Our result captures some of the intuition on how and why panopticism would work in a formal, game-theoretical model.

Moreover, modern society has at its center the individual, not the family or tribe or any other unit. This is crucial for maintaining the self-disciplining aspect of the panopticon, which relies on every prisoner reasoning on his own and choosing what is optimal for him. Others (e.g. Zuboff, 1988) have suggested that modern computers and indeed the internet are panoptica, where everyone can at any time be under surveillance – an idea that has gained credence by recent revelations of mass surveillance by intelligence agencies. Our results, especially the comparison of information structures 2 and 3, suggest that if the true level of surveillance is revealed (or there is a danger of revelation), efficacious enforcement becomes much more expensive in equilibrium – a reason why whistleblowers might indeed pose a threat to enforcement by panopticon. These results show that “order” as used by Foucault, or the central prison metaphor of our theory, are neutral concepts: The free, democratic society might defend itself against an uprising for the sake of social welfare, while a repressive dictatorship might deploy secret surveillance methods to suppress dissent and rebellion. We are interested in the mechanisms by which this is done, and our results are positive, not normative.

Our results also have much more direct applications to situations where one actor can use the coordination problem of his opponents against them. In section 4.1, we discuss the problem of a central bank defending a currency peg against speculators – a question that has

³“Hence the major effect of the Panopticon: to induce in the inmate a state of conscious and permanent visibility that assures the automatic functioning of power. So to arrange things ... that the perfection of power should tend to render its actual exercise unnecessary, ... that the inmates should be caught up in a power situation of which they are themselves the bearers.” (Foucault, 1975)

received much attention in the economic literature (e.g. Flood and Garber, 1984; Obstfeld, 1986; Morris and Shin, 1998). The central bank can – in our model – endogenously build up a costly foreign exchange reserve and decide which information to release about the size of that reserve. We find that the central bank exploits the coordination problem of large groups optimally by maintaining absolute secrecy about its own strength. Similar applications to bank runs and other coordination problems are possible. Our model is also related to works on coordination and revolution, such as Chassang and Miquel (2010) and Edmond (2013). Edmond develops a model in which a government can release biased information about its (exogenously determined) strength – a question that is complementary to our study.

In addition to the mentioned studies, our paper is related to the game-theoretic literature on global games and common knowledge. In the model where the guard level is known, but not common knowledge, we make use of the seminal results on global games; see Carlsson and van Damme (1993), Morris and Shin (1998) and Morris and Shin (2003) for a survey. The “infection” that occurs among prisoners was already described by Rubinstein (1989). We build on this literature but endogenize the “state of nature” as an active choice of the central player, by adding an extra perturbation to the model.

Chwe (2003) provides a discussion of the panopticon and higher-order knowledge. The panopticon, he argues, creates common knowledge among prisoners of being in the same situation – an idea that is connected to Bentham’s plan of having a chapel above the watchtower in his panopticon. Indeed we find that no asymmetric equilibria exist in our panopticon model, i.e. all the prisoners behave exactly the same in equilibrium.

2. Model

This section describes the general setup common to all three models. Details concerning the information structure that differ across the three models are described in the following section.

First, the warden chooses a guard level $\gamma \in \mathbb{R}_+$. Second, N prisoners decide simultaneously and independently whether to revolt (r) or not revolt (n). All revolting prisoners break out if the number of revolting prisoners is strictly larger than γ . Otherwise, no prisoner breaks out. The payoffs are as follows: Each prisoner values breaking out by $b > 0$. If the prisoner revolts but cannot break out, he bears a cost $-q < 0$. This cost can be interpreted in two ways: It could either represent a punishment for prisoners who unsuccessfully try to escape or it could denote a cost of effort (in the latter case b should be interpreted as the benefit of breaking out net of this effort cost). If a prisoner does not revolt, his utility is 0; see table 2 for a summary of these payoffs.

	breaks out	does not break out
r	b	$-q$
n	0	0

Table 2: Payoff prisoner conditional on breaking out or not

The warden experiences a disutility denoted by $-B$ whenever a breakout occurs; apart from that he only cares about the costs of the guards. The costs of the guards are linear in γ with slope normalized to 1, i.e. guard costs are $-\gamma$. Consequently, the utility of the warden is $-B - \gamma$ if a breakout occurs and $-\gamma$ otherwise. Prisoners' and warden's payoffs are assumed to be additive in their components and every player maximizes his expected utility. Finally, we make an assumption on the size of the disutility B . The assumption implies that the warden would prevent a revolt (by setting $\gamma = N$) if he knew that all prisoners play r for sure.

Assumption 1. $B \geq N + 1$.

The reasoning behind this assumption is as follows. If $B < N$, there is – independent of the specific information structure – a very robust equilibrium in which the guard level is zero and all prisoners revolt. This is a somewhat uninteresting case that we want to neglect. For technical reasons, we assume $B \geq N + 1$ (instead of $B > N$) as it significantly simplifies the analysis.

We want to point out two other modeling choices we made: First, the warden's utility depends only on whether there is a breakout and not on how many prisoners break out (or by how much the number of revolting prisoners exceeds the guard level). In this sense, the disutility B corresponds to an image or reputation concern, or a regime preference. Also in the other applications mentioned in the introduction this assumption appears reasonable: A central bank will mainly care about whether it was able to hold the announced peg (and less about how many speculators attacked the peg in case of a successful attack), a government about whether it can stay in power or not. Second, prisoners that do not revolt will not break out (or have at least no benefit from doing so). Think of a prisoner sitting calmly in his cell who will not escape even if others do. Again this fits also the example of speculating against a currency peg: If one does not speculate against the peg, one cannot benefit from a successful attack. It should be noted, however, that our model is robust to deviations from this assumption as long as they do not destroy the strategic complementarity which is at the core of our model – see section 4.2 for details.

3. Analysis

3.1. Benchmark model: Perfect coordination

The first model is a benchmark where we assume the coordination problem of the prisoners away. We distinguish two possibilities: First, the prisoners observe the guard level set by the warden before they have to choose their actions. Assuming the coordination problem away means here that – given the guard level – the prisoners can coordinate on the prisoner optimal Nash equilibrium. Hence, all prisoners play r if $\gamma < N$ and all play n otherwise. Given assumption 1, it is then optimal for the warden to choose $\gamma = N$. The payoff of the warden is $-N$ while the payoff of each prisoner is zero.

Second, we consider the possibility that the prisoners do not observe the guard level. As we allow perfect coordination between the prisoners, prisoners will either all revolt or all not revolt. This is due to the strategic complementarity between prisoners: Revolting is relatively better for a given prisoner if other prisoners revolt too. Given that either all or no prisoners revolt, the only two guard levels that can be best responses by the warden are zero and N . Furthermore, the game has no pure strategy equilibrium because of the non-observability of the guard level: If the warden chose a guard level of zero (N), the prisoners would best respond by revolting (not revolting). But then the guard level of zero (N) is not a best response. Therefore, we only have a mixed equilibrium in which the warden mixes between the two guard levels of zero and N and the prisoners mix between “all revolt” or “no one revolts”. The mixing probabilities are such to keep the other side indifferent. Note that the expected warden payoff is $-N$ since the warden is indifferent between the equilibrium strategy and choosing a guard level of N for sure (which guarantees a payoff of $-N$). The prisoners have an expected payoff of zero as they are indifferent between their equilibrium strategy and not revolting for sure which gives every prisoner a payoff of zero.

Both possibilities of our benchmark lead therefore to the same equilibrium payoffs for all players. In this benchmark model, the warden has to use a large amount of resources to prevent a revolt. The reason is that we assumed that the prisoners had no coordination problem. In the following model, we introduce the coordination problem and show how the warden can exploit this problem to his advantage. In terms of prison design, one might view the benchmark model as a prison in which all prisoners are kept in the same room and find it easy to resolve their coordination problem by communicating with each other. In this interpretation, prisoners are – as Bentham suggested – kept separately in the following models and will therefore face a coordination problem.

3.2. The infection model

In the second model, prisoners choose simultaneously and independently whether to revolt or not. If the guard level is weakly above N , it is a dominant action for each prisoner to play n . If the guard level is strictly below 1, it is a dominant action for each prisoner to play r . For guard levels between 1 and N , the optimal choice of a prisoner depends on what the other prisoners choose: If strictly more than $\gamma - 1$ other prisoners revolt, a given prisoner best chooses r himself. It is, however, optimal to choose n if less than $\gamma - 1$ other prisoners revolt. There are two equilibria in the subgames in which $\gamma \in [1, N)$: All prisoners revolt or no prisoner revolts. Consequently, the prisoners face a coordination problem. Following the approach in the global games literature, we select one of the two equilibria by relaxing the assumption that γ is common knowledge among the prisoners. More precisely, we show that introducing an arbitrarily small amount of noise into the actual and observed guard level leads to a unique equilibrium prediction. Figure 1 shows the intuition behind this equilibrium selection through infection.




Figure 1: Infection of beliefs among prisoners: If $\gamma \geq N$, not revolting is a strictly dominant strategy for all prisoners. If $\gamma < 1$, revolting is strictly dominant. If $\gamma \in [1, N)$ and γ is common knowledge, there are two pure equilibria: Everybody revolts or no one revolts. When common knowledge is destroyed by the perturbation, beliefs get infected so that for $\gamma < \theta^*$, n is the unique equilibrium action, and r is the unique equilibrium action for $\gamma \geq \theta^*$.

The perturbation works in the following way: The warden chooses an intended guard level $\tilde{\gamma}$. The true guard level is then drawn from a normal distribution with mean $\tilde{\gamma}$ and variance $\varepsilon' > 0$.⁴ That is, the warden has a “trembling hand”. Each prisoner receives a noisy signal of γ : This signal is drawn from a uniform distribution on $[\gamma - \varepsilon, \gamma + \varepsilon]$ with $\varepsilon > 0$. We are interested in the Bayesian Nash equilibrium of this game as $\varepsilon \rightarrow 0$. In fact, we show that this Bayesian game has generically a unique Bayesian Nash equilibrium as $\varepsilon \rightarrow 0$. Furthermore,

⁴In the context of a prison, one might think here of a normal distribution truncated at zero. The truncation affects neither results nor derivation.

this equilibrium does not depend on $\varepsilon' > 0$. We select this equilibrium in the original game.⁵

Note that this setup eliminates common knowledge of the guard level. A prisoner observing signal θ knows that the true guard level is in $[\theta - \varepsilon, \theta + \varepsilon]$; he knows that each other prisoner knows that $\gamma \in [\theta - 3\varepsilon, \theta + 3\varepsilon]$; he knows that each other prisoner knows that he knows that $\gamma \in [\theta - 5\varepsilon, \theta + 5\varepsilon]$ etc. Higher order beliefs will therefore play a role in determining the equilibrium. This appears to be a natural feature in a coordination game where the driving force of one's choice are exactly the expectations over what others do (which itself is driven by what others believe I do and therefore beliefs over beliefs and beliefs over beliefs over beliefs etc.).

The following lemma contains the main technical result for the Bayesian game.

Lemma 1. *Let $\varepsilon' > 0$. Assume that $bN/(q + b) \notin \mathbb{N}$ and define⁶*

$$\theta^* = \left\lceil \frac{bN}{q + b} \right\rceil.$$

Then for any $\delta > 0$, there exists an $\bar{\varepsilon} > 0$ such that for all $\varepsilon \leq \bar{\varepsilon}$, a player receiving a signal below $\theta^ - \delta$ will play r and a player receiving a signal above $\theta^* + \delta$ will play n .*

The proposition states that for generic parameter values – whenever $bN/(q + b)$ is not an integer – prisoners in the Bayesian game will revolt when they observe a signal below $\theta^* - \delta$ and will not revolt if they observe a signal above $\theta^* + \delta$. In the limit – as the prisoners' observation noise ε approaches zero – δ approaches zero as well. Put differently, prisoners play a cutoff strategy with cutoff value θ^* in the limit: Whenever they receive a signal below the cutoff, they play r and whenever they receive a signal above the cutoff they play n .

Now consider the warden's decision problem (in the limit as $\varepsilon \rightarrow 0$). If the guard level is strictly above θ^* , then all prisoners will receive signals above θ^* and will therefore not revolt. If the guard level is strictly below θ^* , then all prisoners will receive a signal below θ^* and will revolt. Consequently, the optimal guard level for the warden is θ^* (or slightly above and arbitrarily close θ^*). In the limit as $\varepsilon' \rightarrow 0$, the warden can ensure this guard level by simply choosing $\tilde{\gamma} = \theta^*$. This gives us the following outcome for our second model.

Result 1. *The equilibrium outcome selected by the perturbation is the following: The warden chooses a guard level equal to θ^* and every prisoner plays n .*

⁵The reader familiar with the global games literature might wonder why we introduce a “tremble” in the warden's action. The reason is that the parameter which is observed with noise (the guard level γ) is an endogenous choice in our model while the usual global game approach would assume noisy observation of an exogenous parameter chosen randomly by nature. Since γ is a strategic choice (made before the prisoners act), prisoners could infer γ correctly in equilibrium despite the noisy observation if the warden did not “tremble”. Consequently, prisoners would have common knowledge of γ despite the noise.

⁶The ceiling $\lceil x \rceil$ is the lowest integer above x , i.e. $\lceil x \rceil = \min\{n : n \in \mathbb{N} \text{ and } n > x\}$.

Clearly, the warden does better in this equilibrium than in the benchmark model: He prevents a revolt for sure while using guard level θ^* instead of the guard level N . The reason is that he can utilize the coordination problem among prisoners in his favor. More technically, the so-called “infection argument” is at work: Consider a prisoner receiving a noisy signal above N . It is then quite likely that the guard level is above N and also quite likely that one other prisoner receives a signal above $N + \varepsilon$ (where it is a dominant action to play n). Consequently, a prisoner receiving a signal above N finds it optimal to not revolt. Now consider a prisoner receiving a signal just below N : This prisoner will consider it quite likely that at least one other prisoner receives a signal above N in which case this prisoner will play n (as we just established). So, even if the guard level is below N , it is unlikely that all other prisoners revolt and therefore a prisoner receiving a signal just below N will still play n . In this way, the dominance region (signals above $N + \varepsilon$) “infects” lower and lower signals in the sense that players with these lower signals also find it optimal to play n . A similar infection starts from signals below 1 where it is optimal to play r . Eventually (in the limit), this infection from both sides leads to the unique equilibrium.

3.3. The Panopticon

The third model is the one that comes closest to Bentham’s original idea. Now the warden chooses γ , but it cannot be observed by the prisoners, who also face a coordination problem.⁷⁸ We concentrate on equilibria in which all prisoners play r with the same probability p in equilibrium. In the supplementary material, we show that this is without loss of generality, i.e. no prisoner asymmetric equilibria exist in this game.

Equilibria only exist in mixed strategies: If the prisoners revolted for sure, the warden would best respond by setting the guard level to $\gamma = N$. Consequently, the revolt is unsuccessful and revolting is not a best response for the prisoners. Alternatively, the warden would best respond with $\gamma = 0$ if the prisoners played n for sure. But in this case revolting is a best response. Consequently, the prisoners (and possibly also the warden) will mix and revolts will succeed with some probability in equilibrium.

Since every prisoner plays r with probability p and the prisoners’ choices are independent, the warden faces a binomial distribution of the number of prisoners playing r . Call this

⁷Bentham (1787) emphasized the lack of communication possibilities (leading directly to a coordination problem): “These cells are divided from one another, and the prisoners by that means secluded from all communication with each other, by partitions in the form of radii issuing from the circumference towards the center, and extending as many feet as shall be thought necessary to form the largest dimension of the cell.”

⁸If we allowed prisoners to communicate in a cheap talk way and selected the prisoner optimal equilibrium in this communication game, we would be back in the benchmark model. Such communication is usually not considered in stag hunt type coordination problems as every prisoner weakly benefits if the other prisoner plays revolt; messages are therefore not very credible.

distribution G and its probability mass function g . More precisely, $g(m) = \binom{N}{m} p^m (1-p)^{N-m}$ is the probability that m prisoners revolt given that each prisoner revolts with probability p .

Clearly, the warden's best response puts positive probability only on integers between 0 and N for γ . Therefore, the warden's maximization problem is

$$\max_{\gamma \in \{0, 1, \dots, N\}} -(1 - G(\gamma))B - \gamma. \quad (1)$$

Denote the warden's (mixed) strategy by the distribution F with probability mass function f . The warden has to be indifferent between any two γ_0 and γ_1 in the support of F which means that the following equation has to hold

$$B(G(\gamma_0) - G(\gamma_1)) = c\gamma_0 - \gamma_1 \quad (2)$$

for any γ_0 and γ_1 in the support of F . Note that G is S-shaped because it is a binomial distribution, i.e. g is first strictly increasing (up to the mode of G) and then strictly decreasing. This property leads – together with assumption 1 – to the following result.

Lemma 2. *In any mixed strategy equilibrium, the support of F consists of at most two elements and these two elements are adjacent, i.e. the warden mixes between γ_1 and $\gamma_1 + 1$ with $\gamma_1 \in \{0, \dots, N-1\}$. For any $\gamma_1 \in \{0, \dots, N-1\}$, there exists a unique $p \in (0, (\gamma_1+1)/N)$ such that γ_1 and $\gamma_1 + 1$ are the two global maxima of the warden's utility.*

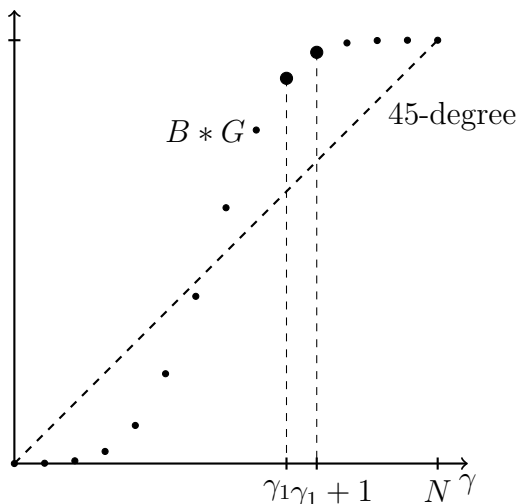


Figure 2: Equilibrium in the panopticon-model.

We illustrate the lemma using figure 2. For every individual revolt probability p , we get a cumulative density function $G(m)$ that gives the probability that m or fewer prisoners

revolt – in other words, the probability that a guard level $\gamma = m$ successfully prevents a breakout. This function G is (multiplied by B) given by the dots (we concentrate on values at integers). The dashed line gives the cost of setting a guard level γ , which is simply γ . The warden optimally mixes between guard levels that maximize the difference between $B * G(\gamma)$ and γ . Intuitively, he trades off the additional cost of increasing the guard strength with reducing the probability of a breakout. Choosing a higher γ than $\gamma_1 + 1$, for example, would increase the cost by much more than the probability of preventing breakouts (weighted by the disutility of a breakout), and is therefore not optimal. If there are several guard levels where the difference is equivalent, the warden is indifferent between them. The example illustrates our two intermediate results: (a) The warden will never mix between more than two guard levels, since the concavity of G (above the mode) means that the difference between G and cost can not be equal in three or more points. (b) For every $\gamma_1, \gamma_1 + 1$ we can find a p such that the warden is indifferent between the two guard levels, by finding a p such that the resulting G has the maximum distance from the 45-degree line at γ_1 and $\gamma_1 + 1$. The condition $p < (\gamma_1 + 1)/N$ is equivalent to saying that the γ_1 is weakly above the mode of G . That is, the optimal guard level will be in the concave part of G which is again in line with figure 2.

In equilibrium, each prisoner must be indifferent between revolting and not revolting. This indifference condition is given by

$$\mathbb{E}_\gamma [-qG_{N-1}(\gamma - 1) + b(1 - G_{N-1}(\gamma - 1))] = 0 \quad (3)$$

where the expectation over γ is taken with respect to the warden's optimal strategy F and G_{N-1} is the binomial distribution with $N - 1$ prisoners, i.e. $g_{N-1}(m) = \binom{N-1}{m} p^m (1-p)^{N-1-m}$. Note that the probability of revolting p and the guard level γ_1 of a mixed equilibrium are determined simultaneously by (1) and (2) as the warden's own mixing probability does not play a role in these conditions. Given these two values, (3) will determine the equilibrium mixing probability of the warden.

We now turn to the question which guard levels can be chosen in equilibrium. Lemma 2 stated that we can concentrate on equilibria where the warden mixes over γ_1 and $\gamma_1 + 1$ for $\gamma_1 \in \{0, \dots, N - 1\}$. Furthermore, the warden's incentives do not pose an obstacle for the existence of such an equilibrium for any $\gamma_1 \in \{0, \dots, N - 1\}$ as there is always a p for which γ_1 and $\gamma_1 + 1$ are optimal. Whether an equilibrium exists for $\gamma_1 \in \{0, \dots, N - 1\}$ is determined by the prisoner's indifference condition. More precisely, a mixed strategy equilibrium where the warden mixes over γ_1 and $\gamma_1 + 1$ exists if and only if the prisoner strictly prefers to revolt if the warden played γ_1 for sure and strictly preferred not to revolt if the warden played $\gamma_1 + 1$

for sure (holding fixed the probability p with which the other prisoners revolt). Defining

$$\Delta(\gamma) = -qG_{N-1}(\gamma - 1) + b(1 - G_{N-1}(\gamma - 1)) \quad (4)$$

as the utility difference of a prisoner between playing revolt and no revolt if the warden uses γ guards for sure, this can be expressed as follows: An equilibrium in which the warden mixes between γ_1 and $\gamma_1 + 1$ exists if and only if $\Delta(\gamma_1) > 0 > \Delta(\gamma_1 + 1)$. In this case, the equilibrium mixing probability with which the warden plays γ_1 is

$$z = \frac{-\Delta(\gamma_1 + 1)}{\Delta(\gamma_1) - \Delta(\gamma_1 + 1)}. \quad (5)$$

Note that several equilibria can exist because Δ is not necessarily monotone: While both terms in (4) are directly decreasing in γ , there is an indirect effect through p : A higher γ is only optimal for the warden if the revolt probability p is higher. This, however, implies that Δ increases. Which of the two effects dominates (direct effect through γ or indirect effect through p) is a priori unclear. However, $\Delta(0) > 0$ as revolting is dominant if the guard level is zero and $\Delta(N) < 0$ as not revolting is dominant when the guard level is N . Consequently, at least one equilibrium exists.

Given that potentially several equilibria exist, we are especially interested in the warden optimal equilibrium. The following lemma shows that the warden optimal equilibrium is the one with the lowest guard level. This equilibrium will also have the lowest revolt probability p .

Lemma 3. *Suppose there are two mixed equilibria: In equilibrium 1, the warden mixes over γ_1 and $\gamma_1 + 1$ and in equilibrium 2 the warden mixes over γ_2 and $\gamma_2 + 1$. Then the warden's equilibrium payoff is higher in equilibrium 1 if and only if $\gamma_1 < \gamma_2$. Furthermore, the prisoners' equilibrium probability of playing r is lower in equilibrium 1 if and only if $\gamma_1 < \gamma_2$.*

So far, we focused on completely mixed equilibria. However, there can be semi-mixed equilibria as well: the warden plays a pure strategy while the prisoners mix. Take a guard level $\gamma \in \{1, \dots, N - 1\}$. There is a range of values for p such that γ is the warden's optimal choice. The prisoner is willing to mix if he is indifferent between revolting and not revolting, that is, if $\Delta(\gamma) = 0$. This indifference condition holds for exactly one p . If the p solving the indifference condition is accidentally within the range of p values for which γ is the maximizer of the warden's utility we have an equilibrium. The following lemma, however, states that semi-mixed equilibria are not warden optimal.

Lemma 4. *For every semi-mixed equilibrium, there is a completely mixed equilibrium in which the expected warden payoff is higher.*

We have therefore established the following for the panopticon model:

Result 2. *In every equilibrium, the prisoners mix over r and n . The warden mixes between some γ_1 and $\gamma_1 + 1$ in the warden optimal equilibrium. However, other equilibria (in which the warden mixes over γ_2 and $\gamma_2 + 1$ with $\gamma_2 > \gamma_1$ or the warden does not mix) can exist.*

3.4. Comparison of the models

The prisoners are indifferent between all models: In the infection model and benchmark 1a, they did not revolt and therefore had a payoff of zero. In the panopticon and benchmark 1b, prisoners were indifferent between revolting and not revolting as they played a mixed strategy. Hence, their expected utility was again zero as this is the payoff from playing n . The warden optimal model will therefore also be the welfare optimal model. Clearly, the benchmark model is worst for the warden: He can prevent a breakout for sure but at very high cost, i.e. his payoff is $-N$. If he prevents communication, he can achieve the same outcome at cost $\theta^* \leq N$. In the panopticon model, he is also weakly better off than in the benchmark, since he always has the option of setting a guard level of N and ensuring a payoff of $-N$. He is indeed indifferent to doing so if the equilibrium in which the warden mixes over $N - 1$ and N is the only existing mixed equilibrium. If other equilibria exist, the warden will, however, be strictly better off in those than in the benchmark model.

The interesting comparison is between the infection model and the panopticon. Which of these two models is warden optimal depends on the parameter values of the model. In general, however, we can show that for large values of N , the panopticon model has a unique equilibrium in which the warden's payoff is bounded from below by a constant. In the infection model, the warden payoff is given by $-\theta^* = -\left[\frac{bN}{q+b}\right]$, which falls linearly in N and therefore becomes very small for large N . We can therefore always find an \bar{N} such that the panopticon is optimal for all $N > \bar{N}$. Also, since the proof of the following proposition establishes that $G(0) \rightarrow 1$ in the unique equilibrium for $N \rightarrow \infty$, the probability of successful breakouts in the panopticon converges to zero.

Proposition 1. *Take b and q as given. Let N be sufficiently large and B such that assumption 1 is satisfied. Then, the warden mixes between 0 and 1 in the unique equilibrium of the panopticon model. The warden's payoff is – for N sufficiently high – higher in this equilibrium than in the infection model.*

To get some intuition for the uniqueness result in the panopticon, consider an equilibrium where the warden mixes over $N - 1$ and N . Assume $B = N + 1$ so that assumption 1 is

satisfied. The warden is only indifferent between the two guard levels if the marginal cost of adding the N th guard, which is 1, equals the marginal benefit of reducing the probability of a breakout by increasing the guard level by one. This marginal benefit is $Bg(N) = (N + 1)p^N$. Hence, $p^N = 1/(N + 1)$ and $p = \sqrt[N]{1/(N + 1)}$ in this equilibrium. Now consider the problem of a prisoner. In this equilibrium, he prefers to revolt only if all other prisoners revolt. The probability that all other prisoners revolt is $g_{N-1}(N - 1) = p^{N-1}$. Since $p = \sqrt[N]{1/(N + 1)}$ by the warden's indifference condition, we get $g_{N-1}(N - 1) = (N + 1)^{-\frac{N-1}{N}}$. This term converges to 0 for large N , so that it becomes extremely unlikely that there is a successful revolt. The prisoner therefore strictly prefers not revolting to revolting, i.e. $\Delta(N - 1) < 0$. Consequently, there is no equilibrium where the warden mixes between N and $N - 1$ for N sufficiently large. A similar logic applies to all other equilibria in which the warden mixes between $\gamma_1 \geq 1$ and $\gamma_1 + 1$: The warden's indifference condition requires a revolt probability p that is – for sufficiently large N – incompatible with the prisoner's indifference condition.

Besides this central result for large groups, we present two results for small N . In this case, either the warden's or the prisoners' payoffs sometimes allow us to say which information structure is optimal.

Proposition 2. *Take q, b, N as given. If $\theta^* = 1$, then the warden is best off in the infection model. If $\theta^* > 1$, then there exists a \bar{B} such that for all $B \geq \bar{B}$ the warden's payoff in the unique equilibrium of the panopticon model is higher than in the infection model. The warden mixes over the guard levels zero and one in this unique equilibrium.*

Put differently, if the disutility of a breakout is relatively high compared to the cost of the guards, the panopticon is warden optimal unless a guard level of 1 can completely deter revolts in the infection model. Given that revolting is dominant for any guard level $\gamma < 1$, $\theta^* = 1$ has to be viewed a bit as a special case. Indeed $\theta^* = \lceil bN/(q + b) \rceil$ equals 1 only if the disutility of an unsuccessful revolt is $N - 1$ times as high as the utility of a successful breakout which seems somewhat implausible in the applications we have in mind. Hence, the panopticon is – with a small caveat – warden optimal if warden incentives dominate. This might be somewhat surprising as the breakout probability in the panopticon is strictly greater than zero while the breakout probability in the infection model is zero. There are two reasons explaining why cost savings compared to the infection model are therefore sizable if $\theta^* > 1$. First, the warden mixes between guard levels of zero and one if B is high. Second, the breakout probability in the panopticon – though not zero – is very small. The second follows readily from the first: Given that the warden really dislikes breakouts (high B), he will only be willing to mix between zero and one if the probability of revolt is very small. The reason why no other equilibrium exists is the following. Given that B is very high, the warden is only willing to use $\gamma_1 < N$ guards if the probability of a revolt is very small. But

this implies that for each prisoner it is unlikely that other prisoners revolt. Consequently, each prisoner strictly prefers not to revolt unless $\gamma_1 = 0$.

Next, consider the prisoners' incentives.

Proposition 3. *Take N and B as given. For b/q high enough, the warden payoff equals $-N$ in all models. Furthermore,*

- *Suppose $B^{\frac{N-1}{N}} > N$: Then, for $b/q \in (N - 1, B^{\frac{N-1}{N}} - 1)$, the warden's payoff in every equilibrium of the panopticon model is higher than in the equilibrium of the infection model.*
- *Suppose $N > B^{\frac{N-1}{N}}$: Then, for $b/q \in (B^{\frac{N-1}{N}} - 1, N - 1)$, there exists an equilibrium in the panopticon model in which the warden's equilibrium payoff is lower than in the infection model.*

If the prisoners have very high incentives to break out, the payoff of all models coincides: The warden chooses N guards in the benchmark 1a and infection model, mixes between N and $N - 1$ guards in the panopticon and between N and 0 in benchmark 1b. Hence, the warden payoff is $-N$. For high (but not excessively high) incentives to break out, the comparison between panopticon and infection model is hampered by the multiplicity of equilibria in the panopticon model. Depending on parameter values, either all (!) equilibria in the panopticon yield a higher warden payoff than the infection model or the infection model does better than some equilibria in the panopticon.

4. Applications and Extensions

4.1. Central Bank Defending Against Speculators

Our results can be applied to many situations of conflict where an agent can use the coordination problem of his opponents against them. An example that has received much attention in economics is the problem of defending a currency peg against speculators. The coordination aspect of this problem, which often leads to multiple equilibria, was pointed out by Flood and Garber (1984) and Obstfeld (1986). The equilibrium multiplicity resulting from the speculators' coordination problem was contentious until Morris and Shin (1998) showed that if speculators lack *common* knowledge about the strength of the currency, their beliefs get infected and there is, for each parameter value, a unique equilibrium. This insight, which builds on the seminal work of Carlsson and van Damme (1993), has since been applied to other coordination problems like bank runs (Goldstein and Pauzner, 2005) or civil war (Chas-sang and Miquel, 2010). These models, however, concentrate on the coordination problem

of the opponents. That is, the underlying “strength” (of the currency, the bank etc.) is exogenous in these models and the setup and information structure is taken as given. Our model, on the other hand, allows us to ask how the agent should defend himself against the coordinated threat and gives direct recommendations as to which information structure is optimal.

Consider the situation of a central bank that has to defend a currency peg against speculation. For this purpose, it can build up foreign exchange holdings that it can then use to counteract speculation. Holding foreign exchange is costly, however, since it requires holding liquid bonds with low yields, so that the central bank would prefer to prevent a breaking of the peg with a minimum of reserves.

The transfer from our prison model is pretty straightforward. Assume that there are N speculators who can each decide to do nothing or take a costly speculative position against the currency. Before the speculators make their choice, the central bank builds up foreign exchange reserves of size γ at cost γ . If there is a speculative attack against the currency and the central bank cannot defend the peg, its payoff is $-B - \gamma$ with some $B \geq N + 1$, otherwise it receives $-\gamma$.

In this context, the assumption that $B \geq N + 1$ means that our model only applies to cases where, if the central bank knew exactly the strength of the speculative attack that was coming, it would always prefer to build a large enough reserve to fight it off. We would argue that this is usually the case in the real world, and that in most cases where a central bank was overwhelmed by speculators it was because of the unexpected extent of the speculative attack.

A speculative attack is successful if more than γ out of the N speculators speculate against the currency. In that case, those who attacked the currency get a payoff of $b > 0$. If they speculate against the currency but the central bank can defend the peg, the speculators lose $q > 0$ on their positions. This loss q denotes the transaction costs of taking the speculating position and also includes the opportunity costs of forgoing an alternative investment. This alternative payoff, which speculators get if they do not speculate against the peg, is normalized to zero.

Now consider the question of whether the central bank should make γ public. Assuming that the bank has the possibility to publicize γ without generating precise common knowledge about it among speculators, this corresponds to a choice between the second and the third information structure in our model.⁹ One way to think about the three information structures is that the bank has the choice between making an official announcement (benchmark model),

⁹If any announcement makes γ common knowledge among speculators, the choice is between the *first* and the third information structure and the choice is clear.

leaking information without officially confirming it (infection model) or giving no information (panopticon).

From our results in the previous sections, we can make several observations about which information policy the central bank should choose in revealing the size γ of the foreign exchange reserve. The optimal choice depends on the interplay of all parameter values, so that the following observations are *ceteris paribus*:

- The central bank should never make a common-knowledge generating announcement about the strength of its currency reserves, for example by using mass media. Such an announcement would instantly solve a coordination problem among speculators and allow them to launch a successful attack. (A common-knowledge generating announcement can only be optimal if the reserve is so strong that it can withstand any attack, in which case the information policy does not matter and the reserve holdings are inefficiently high.)
- If speculators have a lot more to gain from breaking the peg than they can lose by speculating against the peg (in relation to the next-best investment), it may be optimal to keep the reserve level secret. This is especially the case if the cost (economic or reputational) of giving up the peg is high.¹⁰ If the proportion between the speculators' possible earnings and their potential losses grows without bounds, however, the choice of information structure does not matter much since speculators are likely to speculate in any case and the reserve level always has to be maximal.
- If there are many speculators, the central bank should always choose to keep the reserve level secret.

Especially, the third point, which follows directly from our limit result on N , adds a new perspective to the literature on this topic. The uniqueness result of Morris and Shin (1998) has usually been understood to mean that a currency peg can be defended even in cases where coordination among all speculators could bring it down – meaning that a central bank can defend against speculative attacks even if, at this very moment, it doesn't have the power to do so.

But our result shows that while this is true, the central bank can make even better use of the speculators' coordination problem by keeping its own strength secret. Especially if

¹⁰While it may seem like speculators usually have little to lose by speculating against a peg (because they can exchange their money back at the peg if they “lose”), this also includes the cost of transaction and any interest rate differential. Also, re-converting might not be costless if all speculators want to get out at the same time: When the pressure on the Danish krone/Euro peg let off in spring 2015, the Danish central bank suddenly had to stabilize the market *on the other side* of the peg since so many traders reversed or unwound their positions simultaneously.

there are many speculators (i.e. the coordination problem is worse), that will, in the unique equilibrium if N is large enough, guarantee an extremely low probability of losing the peg with a minimal exertion of resources. It should be noted, however, that the massive savings in costly reserves come at the cost of a strictly positive chance of the peg being broken. Observing a central bank that kept its reserves secret being overwhelmed by speculators would, therefore, not necessarily be a sign of a bad policy. While we know of no instance where a central bank actually maintained complete secrecy about the size of its reserves, secrecy about the existence and size of foreign exchange interventions is not uncommon. The reasons for this have been debated in the literature (see Vitale (2007) for a discussion).

4.2. Uncertain punishment: Revolutions, surveillance and prisons

So far, we assumed that revolting leads to a payoff of $-q$ for the prisoner if there was no successful breakout. In particular, this payoff did not depend on the guard level. This is in line with the interpretation of an effort cost in the prison or a transaction cost in the speculation application. One could, however, imagine that revolting prisoners are punished. In the application of a revolution, it is not unreasonable to assume that those that participated in a failed coup d'état might face severe consequences. Punishment, however, requires that the subversive activities are detected. One could argue that the probability of being detected depends on the guard level; i.e. the guards might not detect all unsuccessful revolutionaries if there are few guards monitoring a lot of “prisoners”. One way to capture this is to say that the payoff of a revolting prisoner that does not break out is $-q - \rho\gamma/N < 0$ where $\rho \geq 0$ denotes a punishment and the probability of a punishment is proportional to the guard/prisoner ratio.

As we show in the supplementary material, our analysis covers this more general case. While the specific threshold level θ^* in the infection model and the precise equilibrium mixing probabilities in the panopticon are different, the analysis remains qualitatively the same. In particular, the result that the panopticon is much better than the infection and benchmark model for large N remains true. Also the result that the equilibrium probability of revolting in the panopticon is arbitrarily close to zero for large N holds. This captures one idea sometimes mentioned in connection with the panopticon: The prisoners behave as if they are watched because there is a slight chance that they are watched.¹¹ One could interpret γ/N as the fraction of prisoners that are watched or the chance of being discovered. With $q = 0$ and $\rho > 0$, the only reason not to riot is the possibility of being watched (and punished if caught). Since prisoners almost always do not riot in equilibrium, they arguably behave as if they were watched because they are afraid that they might be watched.

¹¹This idea dates back to Bentham (1787) who writes “You will please to observe, that though perhaps it is the most important point, that the persons to be inspected should always feel themselves as if under inspection, at least as standing a great chance of being so, yet it is not by any means the only one.”

Another possible extension of our model allows the payoff of a non-revolting prisoner to depend on whether a breakout occurs or not. Say, the payoff of a non-revolting prisoner is $w \neq 0$ if a breakout occurs (and zero if no breakout occurs). In the revolution example, w could be negative: If there is a successful coup, the new rulers might punish those that did not participate in the revolt. While the equilibria change quantitatively, all our qualitative results still hold in this setting. The crucial part is that $w < 0$ preserves the supermodular structure of the coordination game: A prisoner is more willing to revolt if other prisoners are more likely to revolt. If, on the other hand, $w > 0$, i.e. if there is a free riding problem, then our results only hold if w is not too big. More precisely, our derivations go through unless the free riding possibility destroys the supermodularity: A prisoner would then be less willing to revolt if others are more likely to revolt because he is more likely to get a high free rider benefit w when not revolting.

5. Conclusion

This paper analyzes how a single player can subdue a group of opponents by making use of their coordination problem. Our model formalizes and replicates earlier results showing that “infection” in the absence of common knowledge can be used for this purpose, but our results go further in arguing that absolute secrecy is often optimal. While secrecy is optimal for all larger groups, the infection model may be optimal for smaller groups of opponents.

In the general debate between secrecy and transparency, this reminds us that we have to think clearly about the purpose and effect of information revelation. Revealing information to a single actor has the effect of informing and influencing that actor, but if that actor is part of a group it will also make him consider what kind of information the others have received, how they reason about his information and so on. These higher-order effects have to be considered and can be substantial.

Our model suggests which is the optimal information structure in a conflict between one central player and a group. However, other situations are conceivable for which our model offers only limited guidance. For example, the idea of transparency and forward guidance by central banks is not necessarily at odds with our result that secrecy is optimal: While our result is based on a conflict between the central bank and speculators, one could imagine other situations in which the interests of central bank and market participants are not opposed. In such a situation with aligned interests, transparency might indeed be an optimal policy. Our results show that the optimal information policy depends crucially on the degree of (mis-)alignment of interests between central bank and market participants.

We have seen that for a large number of prisoners, minimal enforcement with secrecy is

optimal. This is in line with Bentham's original concept. But while prisons indeed rely more on cameras and prisoner separation than on massive numbers of guards, one might wonder why in many other situations massive presence of enforcement is publicly observable. For example, large numbers of police officers are deployed to uphold the public order during (potentially violent) demonstrations and sport events. This is not in contradiction to our theory. Demonstrators (or football hooligans) do not face a large coordination problem. By being in the same place, being able to observe each other and possibly even having some hierarchy among them, they can condition their choices upon each other's behavior and thereby achieve coordination. And, as we have shown in our benchmark model: when coordination problems do not matter, the warden chooses maximum enforcement in equilibrium.

Appendix

Proofs infection model

Proof of lemma 1. The proof is in three steps.

Strategic complementarity: A player finds revolting more attractive if other players are more likely to play revolt. A prisoner's strategy maps from signals into actions.

If there are strategy profiles s and s' such that for every signal for which a player $j \neq i$) plays revolt under s he will also play revolt in s' , then playing revolt is relatively more attractive for player i given s'_{-i} compared to s_{-i} : Define $\Delta(\gamma) = -qG_{N-1}(\gamma - 1) + b(1 - G_{N-1}(\gamma - 1))$ as the utility of revolting minus the utility of not revolting for a given guard level γ . If other players have a higher probability of revolting, then $G_{N-1}(\gamma - 1)$ is weakly lower and therefore $\Delta(\gamma)$ is higher. That is, for a given γ revolting is more attractive. Since this is true for any given γ , it is also true in expectation.

Suppose everyone follows a cutoff strategy with cutoff θ . For a given $\delta > 0$, there exists an $\bar{\varepsilon} > 0$ such that the utility of revolting for a prisoner with signal θ is higher (lower) than the utility from not revolting if $\theta \leq \theta^* - \delta$ ($\theta \geq \theta^* + \delta$). The probability that a player observing himself the cutoff signal θ assigns to the event “exactly k other players receive a signal below θ ” is

$$g_{N-1}(k) = \int_{\theta-\varepsilon}^{\theta+\varepsilon} \binom{N-1}{k} \left(\frac{\gamma - \theta + \varepsilon}{2\varepsilon} \right)^k \left(1 - \frac{\gamma - \theta + \varepsilon}{2\varepsilon} \right)^{N-1-k} \frac{\phi(\gamma)}{\Phi(\theta + \varepsilon) - \Phi(\theta - \varepsilon)} d\gamma.$$

We will now derive a convenient approximation for $g_{N-1}(k)$. Note that for ε small the term $\phi(\gamma)/(\Phi(\theta + \varepsilon) - \Phi(\theta - \varepsilon))$ is approximately constant (and equal to $1/(2\varepsilon)$) as ϕ is continuous and has a bounded first derivative. More precisely, fix θ and define $\phi^{max}(\varepsilon) = \max_{\gamma \in [\theta-\varepsilon, \theta+\varepsilon]} \phi(\gamma)$ and $\phi^{min}(\varepsilon) = \min_{\gamma \in [\theta-\varepsilon, \theta+\varepsilon]} \phi(\gamma)$. Then $g_{N-1}(k)$ and its approximation (where the average $1/(2\varepsilon)$ is used instead of $\phi(\gamma)/(\Phi(\theta + \varepsilon) - \Phi(\theta - \varepsilon))$) are necessarily between the two values (note that the integrand is non-negative for all γ in the integration range)

$$\begin{aligned} \bar{g}(\varepsilon) &= \int_{\theta-\varepsilon}^{\theta+\varepsilon} \binom{N-1}{k} \left(\frac{\gamma - \theta + \varepsilon}{2\varepsilon} \right)^k \left(1 - \frac{\gamma - \theta + \varepsilon}{2\varepsilon} \right)^{N-1-k} \frac{\phi^{max}(\varepsilon)}{\Phi(\theta + \varepsilon) - \Phi(\theta - \varepsilon)} d\gamma, \\ \underline{g}(\varepsilon) &= \int_{\theta-\varepsilon}^{\theta+\varepsilon} \binom{N-1}{k} \left(\frac{\gamma - \theta + \varepsilon}{2\varepsilon} \right)^k \left(1 - \frac{\gamma - \theta + \varepsilon}{2\varepsilon} \right)^{N-1-k} \frac{\phi^{min}(\varepsilon)}{\Phi(\theta + \varepsilon) - \Phi(\theta - \varepsilon)} d\gamma. \end{aligned}$$

By showing that $\lim_{\varepsilon \rightarrow 0} \bar{g}(\varepsilon) - \underline{g}(\varepsilon) = 0$, we show that the approximation of g becomes

arbitrarily close to g for ε small enough:

$$\begin{aligned}\bar{g}(\varepsilon) - \underline{g}(\varepsilon) &= \int_{\theta-\varepsilon}^{\theta+\varepsilon} \binom{N-1}{k} \left(\frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^k \left(1 - \frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^{N-1-k} \frac{\phi^{max}(\varepsilon) - \phi^{min}(\varepsilon)}{\Phi(\theta+\varepsilon) - \Phi(\theta-\varepsilon)} d\gamma \\ &\leq \binom{N-1}{k} \int_{\theta-\varepsilon}^{\theta+\varepsilon} \frac{\phi^{max}(\varepsilon) - \phi^{min}(\varepsilon)}{\Phi(\theta+\varepsilon) - \Phi(\theta-\varepsilon)} d\gamma = \binom{N-1}{k} \frac{2\varepsilon(\phi^{max}(\varepsilon) - \phi^{min}(\varepsilon))}{\Phi(\theta+\varepsilon) - \Phi(\theta-\varepsilon)}.\end{aligned}$$

From L'Hopital's rule and the fact that $\lim_{\varepsilon \rightarrow 0} \phi^{max}(\varepsilon) = \lim_{\varepsilon \rightarrow 0} \phi^{min}(\varepsilon) = \phi(\theta)$, it follows that the last term converges to zero as $\varepsilon \rightarrow 0$. Therefore, the approximation of $g_{N-1}(k)$ converges to $g_{N-1}(k)$ as $\varepsilon \rightarrow 0$. Hence, the approximation is arbitrarily exact for ε sufficiently small (and is totally exact for $\varepsilon = 0$). We will use this result later.

Using the approximation we get

$$\begin{aligned}g_{N-1}(k) &\approx \binom{N-1}{k} \int_{\theta-\varepsilon}^{\theta+\varepsilon} \frac{1}{2\varepsilon} \left(\frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^k \left(1 - \frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^{N-1-k} d\gamma \\ &= \binom{N-1}{k} \int_{\theta-\varepsilon}^{\theta+\varepsilon} \frac{N-1-k}{k+1} \left(\frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^{k+1} \frac{1}{2\varepsilon} \left(1 - \frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^{N-2-k} d\gamma \\ &= \binom{N-1}{k+1} \int_{\theta-\varepsilon}^{\theta+\varepsilon} \left(\frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^{k+1} \frac{1}{2\varepsilon} \left(1 - \frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^{N-2-k} d\gamma\end{aligned}$$

where the step from the first to the second line uses integration by parts (with $[(\gamma - \theta + \varepsilon)/(2\varepsilon)]^k/(2\varepsilon)$ as "first part" and $[1 - (\gamma - \theta + \varepsilon)/(2\varepsilon)]^{N-1-k}$ as "second part"). Using integration by parts for $N - 1 - k$ times gives

$$g_{N-1}(k) \approx \int_{\theta-\varepsilon}^{\theta+\varepsilon} \left(\frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^{N-1} \frac{1}{2\varepsilon} d\gamma = \left[\frac{1}{N} \left(\frac{\gamma-\theta+\varepsilon}{2\varepsilon}\right)^N \right]_{\theta-\varepsilon}^{\theta+\varepsilon} = \frac{1}{N}.$$

Hence, we have obtained that a player receiving the cutoff signal has (approximately) uniform beliefs over the number of players that have received a signal lower than him.

Now we want to consider the expected utility difference between revolting and not revolting of a player receiving cutoff signal θ . If there is no integer $m \in \mathbb{N}$ such that $\theta - \varepsilon \leq m \leq \theta + \varepsilon$, then this utility difference equals $b - (q + b)[\theta]/N$ because a breakout cannot succeed if less than $[\theta]$ other prisoners play revolt.¹² Given the uniform beliefs derived above, the probability that less than $[\theta]$ players play revolt is $[\theta]/N$.

If there is an integer $m \in [\theta - \varepsilon, \theta + \varepsilon]$, then the expected utility difference is

$$b - (q + b) \left[\frac{(\theta + \varepsilon - m)(m + 1)}{2\varepsilon} \frac{1}{N} + \left(1 - \frac{\theta + \varepsilon - m}{2\varepsilon}\right) \frac{m}{N} \right].$$

¹²Recall that $[x] = \max\{n : n \in \mathbb{N} \text{ and } n \leq x\}$, i.e. $[x]$ is the highest integer below x .

Viewed as a function of θ , the expected utility difference is, therefore, flat on intervals (θ_1, θ_2) such that $\lfloor \theta_1 - \varepsilon \rfloor = \lfloor \theta_2 + \varepsilon \rfloor$ and strictly decreasing in an ε -ball around each integer. Hence, there is a unique θ at which the expected utility difference is zero unless the equation $b - (q + b)x/N = 0$ is solved by an integer x , i.e. unless $bN/(q + b) \in \mathbb{N}$, which we ruled out by assumption.¹³ As $bN/(q + b) \in \mathbb{N}$ is clearly not true for generic parameter values (q, b, N) , there exists a unique θ at which the expected utility difference is zero for generic parameter values. In the limit as $\varepsilon = 0$, we then have – for generic parameter values – that (i) the expected utility difference is strictly positive for $\theta < \theta^*$ and (ii) the expected utility difference is strictly negative for $\theta > \theta^*$. Note that (in the limit $\varepsilon = 0$) the expected utility difference (as a function of θ) is discontinuous at θ^* .

The results of the previous paragraph were derived using the approximation of $g_{N-1}(k)$. Now we relax the use of the approximation to obtain the statement we want to show. Take any $\theta < \theta^*$. As the approximation of $g_{N-1}(k)$ converges to $g_{N-1}(k)$, one can find an $\bar{\varepsilon}(\theta) > 0$ such that the expected utility difference is strictly positive for θ for all $\varepsilon \leq \bar{\varepsilon}(\theta)$ (let $\bar{\varepsilon}(\theta)$ be the supremum of all such noise level). Similarly, for each $\theta > \theta^*$ an $\bar{\varepsilon}(\theta)$ can be found such that the expected utility difference at θ is strictly negative for each $\varepsilon \leq \bar{\varepsilon}(\theta)$. Note that $\bar{\varepsilon}(\theta)$ is continuous in θ on $[0, \theta^* - \delta]$ for any given $\delta > 0$: Take $\varepsilon < \bar{\varepsilon}(\theta')$ as given. Since beliefs – i.e. $g_{N-1}(k)$ – change continuously in θ , the expected utility difference is positive not only for θ' but for all θ in some open neighborhood around θ' (given ε). Consequently, $\varepsilon < \bar{\varepsilon}(\theta)$ for every θ in this open neighborhood. A similar argument shows that $\bar{\varepsilon}(\theta)$ is continuous on $[\theta^* + \delta, N]$.

For a given $\delta > 0$, let $\bar{\varepsilon} = \min\{1/2, \min_{\theta \in [0, \theta^* - \delta] \cup [\theta^* + \delta, N]} \bar{\varepsilon}(\theta)\}$. Note that $\min_{\theta \in [0, \theta^* - \delta] \cup [\theta^* + \delta, N]} \bar{\varepsilon}(\theta)$ exists and is strictly greater than zero as it is the minimum over a compact set of an everywhere positive and continuous function. Since revolting is a dominant strategy for signals below $1/2$ (given that $\varepsilon < 1/2$) and not revolting is dominant for signals above $N - 1/2$ (given that $\varepsilon < 1/2$), the expected utility difference is automatically positive (negative) for signals below zero (above N). This concludes the proof of the second step.

For any given $\delta > 0$, there is an $\bar{\varepsilon} > 0$ such that a player with signal below $\theta^* - \delta$ (above $\theta^* + \delta$) plays revolt (not revolt) for all $\varepsilon \leq \bar{\varepsilon}$ in any equilibrium. Hence, each prisoner follows a cutoff strategy with cutoff θ^* in the limit as $\varepsilon \rightarrow 0$. We use the $\bar{\varepsilon}$ determined in step 2. Take an arbitrary equilibrium. Denote by θ_1 the infimum of all signals for which some prisoner does not play revolt for sure. Such a θ_1 exists because of the dominance regions, i.e. revolting (not revolting) is a dominant action for a signal below $1 - \bar{\varepsilon}$ (above $N - 1 + \bar{\varepsilon}$). Then a prisoner receiving any signal below θ_1 should prefer

¹³In this case, the expected utility would be zero on one of the flat parts.

revolting (expected utility difference weakly positive) while there are signals above θ_1 but arbitrarily close to θ_1 where the prisoner prefers not revolting (expected utility difference weakly negative). We will now show that $\theta_1 \geq \theta^* - \delta$: Change all other players strategies such that every player does not revolt if and only if he receives a signal above θ_1 . By the first point, this will make revolting less attractive (decrease the expected utility difference). Hence, a player receiving signal θ_1 will (given that all players use a cutoff strategy with cutoff θ_1) prefer not revolting to revolting. Therefore, by the second step, $\theta_1 \geq \theta^* - \delta$.

Similarly, let θ_2 be the supremum of all signals such that some player plays revolt (with non-zero probability), i.e. for all signals above θ_2 all players prefer not revolting but for some signals below and arbitrary close to θ_2 player i prefers revolting and change the strategies of all other players to cutoff strategies with cutoff θ_2 . Player i will then prefer revolting when receiving signal θ_2 (first step). The second step then implies that $\theta_2 \leq \theta^* + \delta$.

In the limit as $\delta, \varepsilon \rightarrow 0$, we clearly get $\theta_1 = \theta_2 = \theta^*$. □

Proofs and limit results: Panopticon

After the proofs of the results in the main text, we derive another limit result (lemma 5) that we will use when comparing the different models.

Proof of lemma 2. We start with the first part of the lemma. As a first step, we show a weaker result: The support of the warden can consist of at most three elements. Denote the mode of G by γ^m (for a given p).¹⁴ The binomial distribution G has the property that G is convex on $\{0, \dots, \gamma^m\}$ and G is concave on $\{\gamma^m, \dots, N\}$. Therefore, the maximization problem of the warden over the domain $\{0, \dots, \gamma^m\}$ is convex and consequently only the boundary values 0 and γ^m can be local maxima (on this restricted domain). If we take $\{\gamma^m, \dots, N\}$ as domain of the warden's maximization problem, the problem is concave and therefore (because γ takes integer values) this problem can have at most two local maxima γ_1 and γ_2 such that $\gamma_2 = \gamma_1 + 1$ (clearly, it could have only one local maximizer as well in which case we are already done). This implies that (1) has (at most) three local maxima: one at $\gamma_0 = 0$, γ_1 weakly above γ^m and possibly $\gamma_2 = \gamma_1 + 1$. Therefore, f 's support will contain at most three elements.

Next we will show that the case where the warden is indifferent between $\gamma_0 = 0$, $\gamma_1 \geq \gamma^m$ and $\gamma_2 = \gamma_1 + 1$ is impossible. To see this, note that the fact that the warden is indifferent between γ_1 and $\gamma_1 + 1$ implies that $g(\gamma_1 + 1) = 1/B$. The warden is indifferent between γ_1 and γ_0 if and only if $(G(\gamma_1) - G(0))/\gamma = 1/B$. This is equivalent to saying that the average $g(\gamma)$ for $\gamma \in \{1, \dots, \gamma_1\}$ equals $1/B$. Since $\gamma_2 > \gamma^m$ and as $g(\gamma_2) = 1/B$, we know that $g(\gamma) < 1/B$ for all $\gamma > \gamma_2$ (this is true as g is decreasing above the mode). Since $\sum_{\gamma=0}^N g(\gamma) = 1 \geq (N+1)/B$

¹⁴In the non-generic case that G has two modes, let γ^m be the smaller one.

(i.e. the average $g(\gamma)$ is at least $1/B$), this implies that $g(0) \geq 1/B$. But then the single peakedness of g implies that $g(\gamma) > 1/B$ for all $\gamma \in \{1, \dots, \gamma_1\}$ (recall that $g(\gamma_1 + 1) = 1/B$) which contradicts our earlier result that the average $g(\gamma)$ for $\gamma \in \{1, \dots, \gamma_1\}$ is at most $1/B$.¹⁵

Last we reuse the argument of the previous paragraph to show that there cannot be an equilibrium in which the warden mixes between $\gamma_0 = 0$ and $\gamma_1 > 1$. Suppose there was such an equilibrium. Since the warden prefers γ_1 to $\gamma_1 + 1$, we must have $g(\gamma_1 + 1) \leq 1/B$.¹⁶ As γ_1 has to be at least as high as the mode γ^m , we know that $g(\gamma) \leq g(\gamma_1 + 1)$ for all $\gamma \geq \gamma_1 + 1$. The warden prefers γ_1 to $\gamma_1 - 1$ which implies $g(\gamma_1) \geq 1/B$. Furthermore, the warden has to be indifferent between γ_0 and γ_1 which implies that the average $g(\gamma)$ for $\gamma \in \{1, \dots, \gamma_1\}$ equals $1/B$. As $\sum_{\gamma=0}^N g(\gamma) = 1 \geq (N + 1)/B$, we obtain that $g(0) \geq 1/B$. But the single peakedness of g and the fact that $g(\gamma_1) \geq 1/B$ would then imply that the average $g(\gamma)$ for $\gamma \in \{1, \dots, \gamma_1\}$ is strictly above $1/B$ contradicting that the warden is indifferent between γ_0 and γ_1 .

Finally, we turn to the second part of the lemma. Note that $\pi(\gamma_1) = \pi(\gamma_1 + 1)$ holds iff

$$g(\gamma_1 + 1) = 1/B.$$

This equation (viewed as an equation in p which indirectly determines g) has a solution $p < (\gamma_1 + 1)/N$: To see this note that $g(\gamma_1 + 1)$ viewed as a function of p is 0 for $p = 0$ and single peaked with its maximum at $p = (\gamma_1 + 1)/N$. Furthermore, $g(\gamma_1 + 1)$ is continuous in p . Hence, it is sufficient to show that $g(\gamma_1 + 1)|_{p=(\gamma_1+1)/N} > 1/(N + 1)$ as $1/(N + 1) \geq 1/B$ by assumption. Note that for $p = (\gamma_1 + 1)/N$, $\gamma_1 + 1$ is the mode and therefore the maximum of g (viewed as function over γ). If $g(\gamma_1 + 1)|_{p=(\gamma_1+1)/N} \leq 1/(N + 1)$, then $g(\gamma) \leq 1/(N + 1)$ for all γ (with strict inequality for some) which contradicts that g is a probability mass function (it cannot sum to 1!). Hence, $g(\gamma_1 + 1)|_{p=(\gamma_1+1)/N} > 1/(N + 1)$ which proves that there is a $p < (\gamma_1 + 1)/N$ such that $g(\gamma_1 + 1) = 1/B$.

The fact that $p < (\gamma_1 + 1)/N$ implies that $\gamma_1 + 1$ will be above the mode. As π is concave on $\{\gamma^m, \dots, N\}$, it follows that γ_1 and $\gamma_1 + 1$ yield a higher warden payoff than any other γ weakly above the mode. Since π is convex on $\{0, \dots, \gamma^m\}$, it follows that γ_1 and $\gamma_1 + 1$ are global maximizers of π iff $\pi(0) \leq \pi(\gamma_1 + 1)$. This last inequality can be written as

$$\frac{G(\gamma_1 + 1) - G(0)}{\gamma_1 + 1} \geq \frac{1}{B} \tag{6}$$

¹⁵This last argument can be easily extended using inequalities to show that whenever there are γ_1 and $\gamma_2 = \gamma_1 + 1$ forming a local maximum of the warden's profit this local maximum must be the global maximum; i.e. is preferred to $\gamma_0 = 0$.

¹⁶For $\gamma_1 = N$, this step can be skipped and the rest of the argument works analogously.

(where G is the cumulated binomial distribution for the $p < (\gamma_1 + 1)/N$ solving $g(\gamma_1 + 1) = 1/B$). The same argument as above shows that (6) holds: Suppose it did not. Then the average $g(\gamma)$ for $\gamma \in \{1, \dots, \gamma_1 + 1\}$ would be strictly less than $1/B$ and as $\gamma_1 + 1$ is above the mode and $g(\gamma_1 + 1) = 1/B$ the same holds for $\gamma > \gamma_1 + 1$. Using the assumption $B \geq N + 1$ and the fact that $g(\gamma)$ has to sum to 1 over all $\gamma \in \{0, \dots, N\}$, it follows that $g(0) \geq 1/B$. But then the single peakedness of g and $g(\gamma_1 + 1) = 1/B$ contradict that the average $g(\gamma)$ over $\{1, \dots, \gamma_1 + 1\}$ is less than $1/B$. \square

Proof of lemma 3. Let $\gamma_1 < \gamma_2$. We first show that the equilibrium revolting probability p is lower in equilibrium 1. Suppose otherwise, i.e. suppose $p_1 > p_2$. As the warden prefers $\gamma_2 + 1$ over $\gamma_1 + 1$ given p_2 , we have $G^2(\gamma_2 + 1) - G^2(\gamma_1 + 1) \geq (\gamma_2 - \gamma_1)/B$ where G^2 is the binomial cdf under p_2 . This last inequality is equivalent to $\sum_{\gamma=\gamma_1+2}^{\gamma_2+1} g^2(\gamma) - (\gamma_2 - \gamma_1)/B \geq 0$. Note that $\gamma_1 + 1$ is strictly above the mode of g^2 : We know that $\gamma_1 + 1$ is above the mode of g^1 and as $p_1 > p_2$ the mode of g^2 is lower than the mode of g^1 . Similarly, any $\gamma \geq \gamma_1 + 1$ is strictly above the mode of any binomial distribution $g^{(p)}$ with $p \in [p_2, p_1]$. This implies that $\sum_{\gamma=\gamma_1+2}^{\gamma_2+1} g^{(p)}(\gamma) - (\gamma_2 - \gamma_1)/B$ is strictly increasing in p for $p \in [p_2, p_1]$ and therefore $p_1 > p_2$ and $\sum_{\gamma=\gamma_1+2}^{\gamma_2+1} g^2(\gamma) - (\gamma_2 - \gamma_1)/B \geq 0$ imply that $\sum_{\gamma=\gamma_1+2}^{\gamma_2+1} g^1(\gamma) - (\gamma_2 - \gamma_1)/B > 0$. But this is equivalent to saying that the warden strictly prefers $\gamma_2 + 1$ over $\gamma_1 + 1$ under p_1 contradicting that $\gamma_1 + 1$ is the warden's equilibrium choice. Hence, $p_1 > p_2$ cannot hold and we have $p_2 \geq p_1$ whenever $\gamma_2 > \gamma_1$. In fact, $p_2 > p_1$ as otherwise the warden would have to be indifferent between at least three guard levels above the mode which is impossible by the concavity of G on $\{\gamma^m, \dots, N\}$.

Given that $p_2 > p_1$, G^2 first order stochastically dominates G^1 . Therefore, the warden's payoff $-(1 - G(\gamma))B - \gamma$ in equilibrium 1 is higher than his payoff in equilibrium 2 (i.e. if he played γ_2 under p_1 , he would have a higher payoff than in equilibrium 2 and he can do even better by playing γ_1). \square

Proof of lemma 4. Denote by $p(\gamma)$ for $\gamma \in \{0, \dots, N - 1\}$ the value of p for which the warden's payoff is maximized by γ and $\gamma + 1$. The proof of the previous lemma showed that $p(\gamma)$ is strictly increasing in γ . Denote by $\tilde{p}(\gamma)$ the value of p such that $\Delta(\gamma) = 0$.

Now let there be a semi-mixed equilibrium at γ' . This implies that the $\tilde{p}(\gamma')$ is between $p(\gamma' - 1)$ and $p(\gamma')$. If $\tilde{p}(\gamma' - 1)$ is below $p(\gamma' - 1)$, then there is a completely mixed equilibrium where the warden mixes between $\gamma' - 1$ and γ' which leads to a higher payoff for the warden than the γ' equilibrium. Therefore, let's proceed by supposing that $\tilde{p}(\gamma' - 1)$ is above $p(\gamma' - 1)$. This implies that $\tilde{p}(\gamma' - 1)$ is also above $p(\gamma' - 2)$.¹⁷ If $\tilde{p}(\gamma' - 2)$ is below $p(\gamma' - 2)$, then there is a completely mixed equilibrium where the warden mixes between $\gamma' - 1$ and $\gamma' - 2$

¹⁷If $\tilde{p}(\gamma' - 2)$ does not exist, then the prisoner prefers not revolting to revolting for all values of p where $\gamma' - 2$ is weakly above the mode (in particular for $p(\gamma' - 2)$ and $p(\gamma' - 3)$) and the same argument as follows still applies.

which gives him a clearly higher payoff than the γ' equilibrium. Therefore, let us proceed by assuming that $\tilde{p}(\gamma' - 2)$ is above $p(\gamma' - 2)$ which implies that $\tilde{p}(\gamma' - 2)$ is also above $p(\gamma' - 3)$. Iterating further in this way, we finally reach the case where $\tilde{p}(1)$ is above $p(0)$. But this implies that there is an equilibrium where the warden mixes over 0 and 1 and $p = p(0)$: Since $\tilde{p}(1) > p(0)$, $\Delta(1) < 0$ while obviously $\Delta(0) > 0$. \square

Lemma 5. *For sufficiently high b or low q , only the equilibrium in which the warden mixes over N and $N - 1$ exists. For sufficiently high B , the equilibrium in which the warden mixes between 0 and 1 is the only mixed equilibrium.*

Proof. As pointed out in the main text, equilibrium p and γ_1 are determined simultaneously by (2) and (1) as the warden's own mixing probability does not play a role in these conditions. Given these two values, (3) will determine the optimal mixing probability of the warden. This insight shows that b and q will not affect the optimal γ_1 or the equilibrium revolt probability p because these parameters do not play a role in (2) and (1). Note that Δ is linearly increasing in b and linearly decreasing in q . Both variables are not part of the warden's maximization problem. Hence, changes in b and q do not affect the equilibrium mixing probability p for a given support of the warden. This implies that for b high enough (q low enough) $\Delta(\gamma)$ is positive for all $\gamma \in \{0, \dots, N - 1\}$. Hence, only the equilibrium where the warden mixes between $N - 1$ and N exists if b is sufficiently high (or q sufficiently low).

The payoff of the warden when using N guards is $-N$ while his payoff when using $\gamma < N$ guards is $-B(1 - G(\gamma)) - \gamma$. In any mixed equilibrium, the warden has to play an action $\gamma < N$ with positive probability and therefore he must prefer this action (weakly) to the action $\gamma = N$. For $B \rightarrow \infty$, this can only be true if $\lim_{B \rightarrow \infty} p = 0$. Put differently, the equilibrium mixing probability of the prisoner p in a mixed equilibrium becomes arbitrarily small as B increases. Note that very small p imply high $G_{N-1}(\gamma - 1)$ for $\gamma \geq 1$. Consequently, $\Delta(\gamma)$ is negative for sufficiently low p for all $\gamma \geq 1$. As a mixed equilibrium in which the warden mixes over γ_1 and $\gamma_1 + 1$ can only exist if $\Delta(\gamma_1) > 0 > \Delta(\gamma_1 + 1)$, it follows that for sufficiently high B the mixed equilibrium in which the warden mixes over 0 and 1 is the only mixed equilibrium that exists. \square

Proofs model comparison

Proof of proposition 1. It will be convenient to denote $B = \alpha(N+1)$ for some $\alpha \geq 1$ which can be done by assumption 1. In a mixed equilibrium where the warden mixes over 0 and 1, the riot probability p is determined by the warden's indifference condition $1 = BNp(1-p)^{N-1}$. As pointed out in the proof of lemma 2, this p is below $1/N$. The first and main step of in establishing existence of the mixed equilibrium with $\gamma_1 = 0$ (for large N) is to show that

$p < 1/N^2$. By $B = \alpha(N + 1)$ with $\alpha \geq 1$, the indifference condition can be written as $p(1 - p)^{N-1} - 1/(\alpha(N^2 + N)) = 0$. Note that the left hand side of this equation is increasing in p by $p < 1/N$. To show $p < 1/N^2$, it is therefore sufficient to show that the left hand side is greater than 0 for $p = 1/N^2$. This is (after multiplying through by N^2) equivalent to showing that

$$\left(1 - \frac{1}{N^2}\right)^{N-1} > \frac{1}{\alpha\left(1 + \frac{1}{N}\right)}$$

which can be rewritten as

$$\left(1 - \frac{1}{N^2}\right)^N > \frac{1 - 1/N^2}{\alpha\left(1 + \frac{1}{N}\right)} = \frac{N^2 - 1}{\alpha N(N + 1)} = \frac{1 - 1/N}{\alpha}.$$

This inequality holds true as $(1 - 1/N^2)^N = 1 - 1/N + \sum_{i=2}^N \binom{N}{i} (-1/N^2)^i$ and $\sum_{i=2}^N \binom{N}{i} (-1/N^2)^i > 0$ because each positive term in the sum is higher than the immediately following negative term (recall that $\binom{N}{i+1} \leq \binom{N}{i} N$). Given $\alpha \geq 1$, the inequality above therefore holds for all N .

To show that the mixed equilibrium with mixing over 0 and 1 exists, we have to establish that $\Delta(1) < 0$. Given $p < 1/N^2$, $G_{N-1}(0) = (1 - p)^{N-1} > (1 - 1/N^2)^{N-1}$. As $\lim_{N \rightarrow \infty} (1 - 1/N^2)^{N-1} = 1$, this implies that $G_{N-1}(0) \rightarrow 1$ as $N \rightarrow \infty$.¹⁸ Consequently, $\Delta(1) < 0$ for N sufficiently high; i.e. the 0-1 mixed equilibrium exists. Lemma 3 establishes that this is the warden optimal equilibrium in the panopticon.

The warden's payoff in the 0-1 mixed equilibrium is $-B(1 - (1 - p)^N) = -\alpha(N + 1)(1 - (1 - p)^N) > -\alpha(N + 1)(1 - (1 - 1/N^2)^N)$. We now show that the latter term converges to $-\alpha$ as N gets large: This is equivalent to showing that $\lim_{N \rightarrow \infty} N - (N + 1) \left(\frac{N^2 - 1}{N^2}\right)^N = 0$. The term in the limit can be written as

$$\frac{N^{2N+1} - (N + 1)(N^2 - 1)^N}{N^{2N}}.$$

Using the binomial expansion and making use of the fact that $\binom{N}{1} = N$, we can see that this is

$$\frac{N^{2N+1} - N^{2N+1} - N^{2N} + N^{2N} + N^{2N-1} - \dots}{N^{2N}}$$

where the first four terms cancel each other out and the remaining expression only contains powers of N smaller than $2N$ in the numerator, so that the expression goes to zero as N gets large. Therefore, $\lim_{N \rightarrow \infty} (N + 1)(1 - (1 - 1/N^2)^N) = 1$ and the warden's payoff is bounded below by $-\alpha$ in the warden 0-1 mixed equilibrium for N sufficiently large. As the warden's

¹⁸Just to be precise, the limit is 1 as $(1 - 1/N^2)^{N-1} = 1 - N/N^2 + \binom{N}{2} 1/N^4 - \dots$ where all terms but the first approach 0 as N grows large.

payoff is $-\theta^* = -\lceil Nb/(q+b) \rceil$ in the infection model, the warden has a higher payoff in the panopticon for N high enough.¹⁹

Finally, we show uniqueness of the mixed equilibrium with $\gamma_1 = 0$ in the panopticon (for large N). To do so, we need two intermediate results that are stated as lemmas below (lemma 6 and 7). To start with, define an *equilibrium candidate* as a (p, γ) such that the warden's indifference condition holds, that is $g(\gamma + 1) = \frac{1}{\alpha(N+1)}$, and $p < (\gamma + 1)/N$. An equilibrium candidate leads to an equilibrium if $\Delta(\gamma) \geq 0$ and $\Delta(\gamma + 1) < 0$, that is if $G_{N-1}(\gamma - 1) \leq b/(q+b) \leq G_{N-1}(\gamma)$. We will show that for large N , there are no equilibrium candidates with $\gamma \geq 1$ that satisfy the equilibrium condition $G_{N-1}(\gamma - 1) \leq b/(q+b)$.

In the following, we make use of known results on the shape and the tail bounds of the binomial distribution. Recall that $g_N(\gamma) = \binom{N}{\gamma} p^\gamma (1-p)^{N-\gamma}$, i.e. the probability mass of the binomial distribution $B(N, p)$ at γ . G_N is the corresponding cumulative distribution function; the definitions of g_{N-1} and G_{N-1} are analogous.

Lemma 6. *The probability $1 - G_N(\gamma)$ that $\gamma + 1$ or more prisoners revolt in any equilibrium candidate (and therefore the probability of a breakout) converges to zero as N grows large.*

Proof. Using the Chernoff bound (Chernoff, 1952), we get

$$1 - G_N(\gamma) \leq \left(\frac{N}{\gamma + 1}\right)^{\gamma+1} \left(\frac{N}{N - \gamma - 1}\right)^{N-\gamma-1} p^{\gamma+1} (1-p)^{N-\gamma-1}. \quad (7)$$

For any equilibrium candidate in which the warden mixes over γ and $\gamma + 1$, it is therefore

$$1 - G_N(\gamma) \leq \left(\frac{N}{\gamma + 1}\right)^{\gamma+1} \left(\frac{N}{N - \gamma - 1}\right)^{N-\gamma-1} \frac{1}{\alpha(N+1) \binom{N}{\gamma+1}}$$

where we plug the warden's indifference condition into (7). It is convenient to define $m = \gamma + 1$ as this allows to write the previous expression as

$$1 - G_N(\gamma) \leq \frac{N^N}{\binom{N}{m} m^m (N-m)^{N-m} \alpha(N+1)}. \quad (8)$$

We are going to show that the RHS term converges to zero as N grows large. We have to show this for any $m \in \{1, \dots, N\}$ and in particular m might depend on N . That is, we want to show that the expression above converges to zero for any $m(N)$. To do so, let $m^*(N)$ be

¹⁹Note that the result does not depend on using a fixed α . More precisely, take a sequence of N and $B_N = \alpha_N(N+1)$ with $\alpha_N \geq 1$ for all N . The previous steps above still apply (for each given N) and the warden will prefer the no information 0-1 mixed equilibrium to $-\theta^*$ for N high enough as long as the sequence of α_N is bounded by some $\bar{\alpha}$.

the m maximizing the expression above. We show that the expression converges to zero even if we plug in $m = m^*(N)$.

Note that the term in (8) is maximal (for a given N) if m minimizes $\binom{N}{m}(m/N)^m(1 - m/N)^{N-m}$. Note that $\binom{N}{m}(m/N)^m(1 - m/N)^{N-m}$ is the probability mass of a binomial distribution with probability $p = m/N$ evaluated at its mode m . Hence, to minimize $\binom{N}{m}(m/N)^m(1 - m/N)^{N-m}$ we have to find the probability $p = m/N$ for which the modal density of a binomial distribution is minimized. This is the case for $p = 1/2$, i.e. $m = N/2$.²⁰ Consequently, $\forall m(N) : \binom{N}{m}m^m(N - m)^{N-m} \leq \left(\frac{N}{2}\right) \left(\frac{N}{2}\right)^N$ and (8) becomes

$$\begin{aligned} 1 - G_N(\gamma) &\leq \frac{N^N}{\binom{N}{N/2}(N/2)^N \alpha(N+1)} \\ &= \frac{2^N}{\binom{N}{N/2} \alpha(N+1)}. \end{aligned} \tag{9}$$

Since the central binomial coefficient $\binom{N}{N/2}$ is bounded from below by $2^N/\sqrt{2N}$ (see the supplementary material for an elementary proof of this), we obtain that $1 - G(\gamma)$ converges to zero in any equilibrium candidate. \square

We will now use this result to show that not only the probability of successful revolts converges to zero, but also the probability for each prisoner that a revolt will be successful if he decides to revolt. This is given by $1 - G_{N-1}(\gamma - 1)$, i.e. the probability that at least γ other prisoners revolt (so that the remaining prisoner can push the number to $\gamma + 1$ or higher by revolting himself).

Lemma 7. *In any equilibrium candidate with $\gamma \geq 1$, $1 - G_{N-1}(\gamma - 1)$ converges to zero as N grows large.*

Proof. Note that $1 - G_{N-1}(\gamma - 1) = 1 - G_{N-1}(\gamma) + g_{N-1}(\gamma) \leq 1 - G(\gamma) + g_{N-1}(\gamma)$. From lemma 6 we know that $1 - G(\gamma)$ converges to zero in any equilibrium candidate. It is therefore sufficient to show that $g_{N-1}(\gamma)$ converges to zero in any equilibrium candidate as N grows large. We distinguish two cases. First, $g_{N-1}(\gamma)$ converges to zero as N grows large. In this case, we are done. Second, $g_{N-1}(\gamma)$ does not converge to zero. We will show directly that $1 - G_{N-1}(\gamma - 1)$ converges to zero for large enough N in this case.

By the warden's indifference condition, $g_N(\gamma + 1) = \frac{1}{\alpha(N+1)}$, and we can write

$$g_{N-1}(\gamma) = g_N(\gamma + 1) \frac{\gamma + 1}{pN} = \frac{\gamma + 1}{\alpha p N^2}.$$

²⁰If N is odd, both $m = \lfloor N/2 \rfloor$ and $m = \lceil N/2 \rceil$ will lead to minimal modal density. We concentrate on the case where N is even for notational convenience. Obviously, our results also hold for odd N .

If this does not converge to zero, there is a sequence of tuples (N, p, γ) which is strictly increasing in N such that (i) (p, γ) is an equilibrium candidate for each tuple (N, p, γ) and (ii) $\gamma + 1 \geq \mu p N^2$ for each tuple in the sequence and some $\mu > 0$.

Rearranging the latter condition gives

$$\gamma - pN + p \geq \mu p N^2 - pN + p - 1 = pN^{5/4} \left(\mu N^{3/4} - \frac{1}{N^{1/4}} \right) + p - 1. \quad (10)$$

We will look at two cases. First, $pN^{5/4}$ does not converge to zero. Then the right hand side of (10) is weakly larger than $\tilde{\mu} N^{3/4}$ for some $\tilde{\mu} > 0$ and N sufficiently large. Therefore, $\frac{(\gamma - pN + p)^2}{N-1} \geq \frac{(\tilde{\mu} N^{3/4})^2}{N-1} > \tilde{\mu}^2 \sqrt{N}$ for large N which implies that $\frac{(\gamma - pN + p)^2}{N-1}$ will grow without bound as N gets large. Hoeffding's inequality (Hoeffding, 1963, Thm. 1) gives the following upper bound for $1 - G_{N-1}(\gamma - 1)$:

$$1 - G_{N-1}(\gamma - 1) \leq e^{-\frac{2(\gamma - p(N-1))^2}{N-1}}.$$

As we have just shown, this upper bound tends to zero as N grows large. Consequently, we have shown directly that $1 - G_{N-1}(\gamma - 1)$ converges to zero. It remains to check the second case in which $pN^{5/4}$ converges to zero. If $pN^{5/4}$ converges to zero, then $p \leq 1/N^{5/4}$ for sufficiently high N . Consequently, $G_{N-1}(0) = (1 - p)^N \geq (1 - 1/N^{5/4})^N$ and the latter converges to 1. As $G_{N-1}(0) \leq G_{N-1}(\gamma - 1)$ for $\gamma \geq 1$, this implies that $1 - G_{N-1}(\gamma - 1)$ converges to zero. \square

Lemma 7 implies that $G_{N-1}(\gamma - 1)$ converges to one for any equilibrium candidate with $\gamma \geq 1$ as N gets large. Put differently, for any $\varepsilon > 0$, we can find an $\bar{N}(\varepsilon)$ such that $G_{N-1}(\gamma_1) > 1 - \varepsilon$ for all $N \geq \bar{N}(\varepsilon)$ and all equilibrium candidates with $\gamma \geq 1$. In particular, we can find such an $\bar{N}(\varepsilon)$ for $\varepsilon = 1 - b/(q + b)$. For $N \geq \bar{N}(1 - b/(q + b))$, we have $G_{N-1}(\gamma - 1) > b/(q + b)$ for all equilibrium candidates with $\gamma \geq 1$. Hence, no equilibrium candidate with $\gamma \geq 1$ satisfies the equilibrium condition $G_{N-1}(\gamma - 1) \leq b/(q + b)$ for N sufficiently high. This means that the equilibrium in which the warden mixes over zero and one is the unique equilibrium for N sufficiently high. \square

Proof of proposition 2. Lemma 5 establishes that for B high enough the only mixed equilibrium is the one where the warden mixes over 0 and 1. The proof of the lemma also establishes that $\Delta(\gamma) < 0$ for $\gamma \geq 1$ if B is sufficiently high. Consequently, also no semi-mixed equilibrium exists for B high enough. Let \hat{B} be such that only the mixed equilibrium in which the warden mixes over 0 and 1 exists for any $B \geq \hat{B}$. For the rest of the proof, consider only $B \geq \hat{B}$.

In this mixed equilibrium the warden is indifferent between 0 and 1 which means $Bg(1) = 1$ or equivalently $N(1 - p)^{N-1}p = 1/B$. Therefore, $\lim_{B \rightarrow \infty} p(B) = 0$ where $p(B)$ is the

prisoners' equilibrium probability of playing r when the warden's utility is B . Since the warden is indifferent between playing 0 and 1 in equilibrium, his equilibrium payoff equals $\pi(0) = -(1 - (1 - p)^N)B$. Plugging in the indifference condition $N(1 - p)^{N-1}p = 1/B$ derived above yields the warden's equilibrium payoff

$$\pi^* = \frac{(1 - p)^N - 1}{N(1 - p)^{N-1}p}.$$

Applying L'Hôpital's rule, gives $\lim_{p \rightarrow 0} \pi^* = -1$. As we established above, p approaches 0 when $B \rightarrow \infty$. Consequently, the warden's payoff in the mixed equilibrium approaches -1 as $B \rightarrow \infty$. Furthermore,

$$\begin{aligned} \frac{\partial \pi^*}{\partial p} &= \frac{-N^2(1 - p)^{2N-2}p - ((1 - p)^N - 1)(-N(N - 1)(1 - p)^{N-2}p + N(1 - p)^{N-1})}{N^2(1 - p)^{2N-2}p^2} \\ &= \frac{1 - Np - (1 - p)^N}{N(1 - p)^N p^2}. \end{aligned}$$

Using L'Hôpital's rule, gives $\partial \pi^* / \partial p|_{p=0} = -(N - 1)/2 < 0$. Hence, the warden's payoff approaches -1 from below as $B \rightarrow \infty$ and the warden's payoff in the equilibrium where he mixes over 0 and 1 is bounded from above by -1 . This proves the proposition because in the infection model the warden's equilibrium payoff is $-\theta^*$ for any value of B .²¹ \square

Proof of proposition 3. It was shown before that for b/q high enough, the unique equilibrium in the panopticon model is a mixed equilibrium in which the warden mixes over $N - 1$ and N and his payoff is $-N$. A similar result holds for the infection model: $\theta^* = N$ if and only if $b/(q + b) > (N - 1)/N$ or equivalently if $(b/q) > N - 1$. Clearly, $\theta^* = N$ implies that the warden's equilibrium payoff is $-N$. This establishes the result that for b/q high enough all models lead to a warden payoff of $-N$.

Now consider the panopticon. In an equilibrium in which the warden mixes over $N - 1$ and N , he has to be indifferent between these two options which implies $1 = Bp^N$, i.e. the mixing probability of the prisoner has to be $p = (1/B)^{1/N}$ in such an equilibrium. To have such an equilibrium, the condition $\Delta(N - 1) > 0$ has to be satisfied. Given $p = (1/B)^{1/N}$, this condition becomes $-q(1 - (1/B)^{(N-1)/N}) + b(1/B)^{(N-1)/N} > 0$. This can be rewritten as $b/q > (1/B)^{(N-1)/N} - 1$.

If $B^{(N-1)/N} - 1 > b/q > N - 1$, then the warden's payoff in the infection model is $-N$. In the panopticon, however, the equilibrium in which the warden mixes between N and $N - 1$ does not exist which means the warden plays N with zero probability in any equilibrium of

²¹Note that for low values of B where other equilibria might exist the warden's payoff is still bound from below by -1 : In any such equilibrium the warden finds it optimal to use a guard level of 1 or higher and the breakout probability is positive. Consequently, the warden's expected payoff is strictly lower than -1 .

this game. As the equilibrium guard levels are then strictly preferred to a guard level of N (which would guarantee payoff $-N$), it follows that the warden's payoff in the no information game is strictly larger than $-N$.

If $B^{(N-1)/N} - 1 < b/q < N - 1$, the no information game has an equilibrium in which the warden mixes between $N - 1$ and N and therefore his expected payoff in this equilibrium is $-N$. In the infection game, $\theta^* < N$ and therefore the warden's equilibrium payoff is strictly above $-N$. □

References

- Bentham, J. (1787). *Panopticon; Or, The Inspection-House*. The Works of Jeremy Bentham, published under the superintendence of his executor John Bowring (Edinburgh: William Tait, 1838-1843). 11 vols. Vol. 4.
- Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica* 61(5), 989–1018.
- Chassang, S. and G. P. I. Miquel (2010). Conflict and deterrence under strategic risk. *Quarterly Journal of Economics* 125(4), 1821–1858.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics* 23(4), 493–507.
- Chwe, M. S.-Y. (2003). *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton: Princeton University Press.
- Edmond, C. (2013). Information manipulation, coordination, and regime change. *Review of Economic Studies* 80, 1422–1458.
- Flood, R. P. and P. M. Garber (1984). Collapsing exchange rate regimes: Some linear examples. *Journal of International Economics* 17, 1–13.
- Foucault, M. (1975). *Discipline and Punish: The Birth of the Prison* (trans. Alan Sheridan). New York: Vintage Books.
- Goldstein, I. and A. Pauzner (2005). Demand-deposit contracts and the probability of bank runs. *Journal of Finance* 60(3), 1293–1327.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301), 13–30.
- Morris, S. and H. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88(3), 587–597.
- Morris, S. and H. S. Shin (2003). Global games: theory and applications. In M. Dewatripont, L. Hansen, and S. Turnovsky (Eds.), *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*. Cambridge: Cambridge University Press.
- Obstfeld, M. (1986). Rational and self-fulfilling balance-of-payments crises. *American Economic Review* 76(1), pp. 72–81.

- Rubinstein, A. (1989). The electronic mail game: Strategic behavior under almost common knowledge. *American Economic Review* 79(3), 385–391.
- Vitale, P. (2007). An assessment of some open issues in the analysis of foreign exchange intervention. *International Journal of Finance and Economics* 12, 155–170.
- Zuboff, S. (1988). *In the age of the smart machine: the future of work and power*. New York: Basic Books.

Supplementary Material

not intended for publication

Panopticon: No asymmetric equilibria

When analyzing the panopticon model, we restricted attention to symmetric equilibria, i.e. equilibria in which all prisoners revolt with the same probability p . We will now show that this is without loss of generality, i.e. there are no equilibria in which prisoners revolt with prisoner dependent probabilities p_i and $p_i \neq p_j$ for some prisoners i and j .

In the main text, we already argued that equilibria cannot be pure, i.e. there has to be at least one prisoner who uses a mixed strategy p_i with $0 < p_i < 1$. The argument is simple: If all prisoners used a pure strategy in equilibrium, the warden would be certain of the number of revolting prisoners, say k . In this case, the warden best responds by setting $\gamma = k$ which prevents a breakout for sure while any lower guard level would lead to a breakout with probability 1. If $k > 0$, the revolting prisoners could profitably deviate to not revolting. If, however, $\gamma = k = 0$, then each prisoner could profitably deviate by revolting. Since at least one prisoner has a profitable deviation, we can conclude that there is no equilibrium in which all prisoners use pure strategies. Without loss of generality, let us therefore assume that prisoner 1 uses a completely mixed strategy, i.e. $0 < p_1 < 1$.

First, we will show the following: Take any equilibrium in the panopticon model. If $0 < p_i \leq p_j < 1$ holds for two prisoners i and j , then $p_i = p_j$. To see this, note that both i and j have to be indifferent between revolting and not revolting because both use a completely mixed strategy. If $p_j > p_i$ and j is indifferent between revolting and not revolting, then i would strictly prefer to revolt: For any $\gamma > 0$, the probability that at least $\lfloor \gamma \rfloor$ other prisoners revolt is higher for i than for j if $p_j > p_i$. Since j was indifferent, i will then strictly prefer to revolt. This contradicts that i is indifferent (because he plays a completely mixed strategy) and we must therefore have $p_i = p_j$.

Note that the previous argument actually says that if two players are indifferent between revolting and not revolting, then they must play revolt with the same probability. This is a bit stronger than what we said before because it rules out the possibility that some prisoner plays revolt with probability 0 or 1 while being indifferent between the two actions. (Recall that prisoner 1 uses a completely mixed strategy.)

What remains to be shown is that no prisoner strictly prefers one of the two actions in equilibrium. Suppose to the contrary that prisoner j strictly preferred to revolt and therefore plays revolt with probability 1 in equilibrium. Now consider prisoner 1: Since $p_1 < p_j = 1$, the probability that at least $\lfloor \gamma \rfloor$ other prisoners revolt is higher from prisoner 1's perspective

than from prisoner j 's perspective. Therefore, prisoner 1 strictly prefers to revolt given that prisoner j strictly prefers to revolt. This contradicts that prisoner 1 plays a completely mixed strategy in equilibrium. Consequently, there cannot be a prisoner j who strictly prefers to revolt.

An analogous argument yields that there is no prisoner who strictly prefers not revolt. This completes the proof.

Extension: Uncertain punishment

Here we consider a variation of the model in which a prisoner's payoff when revolting unsuccessfully is $-q - \rho\gamma/N < 0$ where $q \geq 0$ is an effort cost and $\rho \geq 0$ is a punishment that happens with probability γ/N . It will become apparent that the the specific linear form chosen here is irrelevant for the analysis, i.e. we could just as well use $-q - h(\gamma, N)$ where $h \geq 0$ increases in its first and decreases in its second argument. Apart from this change in payoff, the model is the same as in the main text.

Note that the arguments in the **benchmark model** go through without change.

In the **infection model**, lemma 1 holds with a slightly redefined threshold θ^* . Let θ^* be the unique θ such that

- either $\theta \notin \mathbb{N}$ and

$$b - \left(q + b + \frac{\theta}{N}\rho \right) \frac{\lfloor \theta \rfloor}{N}$$

- or $\theta \in \mathbb{N}$ and

$$\begin{aligned} 0 &\geq b - \left(q + b + \frac{\theta}{N}\rho \right) \frac{\theta}{N} \\ 0 &\leq b - \left(q + b + \frac{\theta}{N}\rho \right) \frac{\theta - 1}{N}. \end{aligned}$$

The proof of lemma 1 has to be adjusted only at very few instances: In the first step,

$$\Delta(\gamma) = b - \left(q + b + \frac{\theta}{N}\rho \right) G_{N-1}(\gamma - 1)$$

and everything goes through accordingly.

In the second step, the derivation of the approximation and the resulting Laplacian beliefs remains unaffected. The expected utility difference between rioting and not rioting if there

does not exist an $m \in \mathbb{N}$ such that $\theta - \varepsilon \leq m \leq \theta + \varepsilon$ will now be

$$b - \left(q + b + \frac{\theta}{N} \rho \right) \frac{\lfloor \theta \rfloor}{N}.$$

If such an m exists, the expected utility difference is

$$b - \left(q + b + \left(\frac{m}{2} + \frac{\theta + \varepsilon}{2} \right) \frac{\rho}{N} \right) \frac{\theta + \varepsilon - m}{2\varepsilon} \frac{m + 1}{N} - \left(q + b + \left(\frac{m}{2} + \frac{\theta - \varepsilon}{2} \right) \frac{\rho}{N} \right) \left(1 - \frac{\theta + \varepsilon - m}{2\varepsilon} \right) \frac{m}{N}.$$

Note that this expected utility difference is strictly decreasing in θ if $\rho > 0$. As rioting is dominant for $\theta < 1 - \varepsilon$ and not rioting is dominant for $\theta > N + \varepsilon$, there is a unique θ at which the expected utility difference is zero. In the limit $\varepsilon \rightarrow 0$, we obtain that the expected utility difference is strictly positive for every $\theta < \theta^*$ and strictly negative for every $\theta > \theta^*$. Given this, the remaining parts of the proof of lemma 1 apply without change.

In the **panopticon model**, the indifference condition of the prisoner (3) has to be rewritten as

$$\mathbb{E} \left[b - G_{N-1}(\gamma - 1) \left(b + q + \rho \frac{\gamma}{N} \right) \right] = 0.$$

Lemmas 2 and 3 remain valid because they use only the warden's problem which was not changed. The proofs of lemmas 5 and 4 use the prisoners' indifference condition without using the specific form of the prisoner payoff. Consequently, the proofs go through without change and the lemmas remain valid.

The most interesting **comparison** of the models is the result for large N (proposition 1). The proof of this result does again not use the specific form of the prisoners' indifference condition and consequently goes through without change. Hence, all the results for large N mentioned in the main text remain valid.

Example: N=2

To illustrate the results of the paper, we give the solved model for the simple case where $N = 2$.

Denoting the expected warden payoff by $\pi(\gamma)$, we get for the $N = 2$ case

$$\begin{aligned} \pi(0) &= -(2p + p^2)B \\ \pi(1) &= -p^2B - 1 \\ \pi(2) &= -2. \end{aligned}$$

This implies that $\pi(0) = \pi(1)$ iff $p = 1/(2B)$. Given the assumption $B \geq N + 1 = 3$,

$\pi(0) = \pi(1) > \pi(2)$ holds if $p = 1/(2B)$.

Furthermore, $\pi(1) = \pi(2)$ iff $p = \sqrt{\frac{1}{B}}$ and $B \geq 3$ implies in this case that $\pi(1) = \pi(2) > \pi(0)$. To determine the equilibrium we will have to check the prisoners' indifference condition. Denoting the utility difference from revolting and not revolting given γ guards by $\Delta(\gamma)$ we get

$$\begin{aligned}\Delta(0) &= b \\ \Delta(1) &= -q(1-p) + bp \\ \Delta(2) &= -q.\end{aligned}$$

If $\Delta(1) < 0$ with $p = 1/(2B)$, then there is an equilibrium in which the warden mixes over 0 and 1 with probability $z_{0,1} = \frac{-\Delta(1)}{-\Delta(1)+\Delta(0)} = \frac{q-b/(2B-1)}{q+b}$. The inequality $\Delta(1) < 0$ is, given $p = 1/(2B)$, equivalent to $b/q < 2B - 1$.

If $\Delta(1) > 0$ with $p = \sqrt{\frac{1}{B}}$, then there exists an equilibrium in which the warden mixes over 1 and 2 with probability $z_{1,2} = \frac{q}{p(b+q)} = \sqrt{B} \frac{q}{q+b}$. Then the inequality $\Delta(1) > 0$, given $p = \sqrt{1/B}$, is $b/q > \sqrt{B} - 1$.

Note that $\sqrt{\frac{1}{B}} > 1/(2B)$ and $2B - 1 > \sqrt{B} - 1$ by $B \geq N + 1 = 3$. This implies the structure in figure 3 for existence of the different equilibria.




Figure 3: Equilibria for N=2 case

The warden payoff in the 0,1 mixing equilibrium equals $\pi(1) = -p^2B - 1 = -\frac{1}{4B} - 1$. The warden payoff in the 1,2 mixing equilibrium equals $\pi(2) = -2$.

Last, we look at semi-mixed equilibria, i.e. the warden plays a pure strategy while the prisoners play completely mixed strategies. Note that the warden cannot play the pure strategies 0 or 2 in such an equilibrium because the prisoners would then have a dominant action contradicting that they mix. Hence, we can focus on the equilibrium where the warden plays $\gamma = 1$. Playing $\gamma = 1$ is optimal for the warden if $p \in \left[1/(2B), \sqrt{1/B}\right]$. The prisoner is willing to mix only if $\Delta(1) = 0$, i.e. if $b/q = (1-p)/p = 1/p - 1$. Note that $1/p - 1$ equals $2B - 1$ for $p = 1/(2B)$ and $1/p - 1$ equals $\sqrt{B} - 1$ for $p = \sqrt{1/B}$. Consequently, the

semi-mixed equilibrium exists if $\frac{b}{q} \in [\sqrt{B} - 1, 2B - 1]$.

The warden payoff in the panopticon were already established above. In particular, the mixed equilibrium with mixing over zero and one existed if $b/q < 2B - 1$ and the warden payoff in this game was $-1/(4B) - 1$. For $b/q > 2B - 1$, only the mixed equilibrium with mixing over 1 and 2 existed where the warden payoff is -2. In the infection model, $\theta^* = 1$ if $b/q < 1$ and $\theta^* = 2$ if $b/q > 1$. This implies that the warden payoff is higher in the infection model than in the panopticon if $b/q < 1$. For $1 < b/q < 2B - 1$, the warden optimal equilibrium of the panopticon gives the warden a higher payoff than the infection model. The worst equilibrium in the panopticon model gives the warden the same payoff as the infection model in this case. If $b/q > 2B - 1$, all models give payoff -2 to the warden.

Lower bound of the central binomial coefficient – Proof²²

We will show the equivalent $\binom{2n}{n} \geq 2^{2n}/(2\sqrt{n})$ as it is notationally more convenient. The first step is to see that

$$\begin{aligned} \binom{2n}{n} \frac{1}{2^{2n}} &= \frac{1}{2^{2n}} \frac{(2n)!}{n!n!} \\ &= \frac{1}{2^n} \frac{(2n)!}{n!2^n n!} \\ &= \frac{1}{2^n} \frac{(2n-1)(2n-3)(2n-5)\dots 1}{n!} \\ &= \frac{1}{2^{n-1}} \frac{1}{2n} \frac{(2n-1)(2n-3)(2n-5)*\dots*3}{(n-1)(n-2)*\dots*1} \\ &= \frac{1}{2^{n-1}} \frac{1}{2n} \prod_{j=1}^{n-1} \frac{2j+1}{j} \\ &= \frac{1}{2n} \prod_{j=1}^{n-1} \left(1 + \frac{1}{2j}\right). \end{aligned}$$

The second step is to get a lower bound on the square of the product:

$$\begin{aligned} \prod_{j=1}^{n-1} \left(1 + \frac{1}{2j}\right)^2 &= \prod_{j=1}^{n-1} \left(1 + \frac{1}{j} + \frac{1}{4j^2}\right) \\ &\geq \prod_{j=1}^{n-1} \left(1 + \frac{1}{j}\right) = n. \end{aligned}$$

²²The proof is a slightly more extensive version of a proof given by Byron Schmuland on <http://math.stackexchange.com/questions/58560/elementary-central-binomial-coefficient-estimates>.

Where the last equality can be easily shown by induction.²³ Taking the first two steps together shows that

$$\left(\binom{2n}{n} \frac{1}{2^{2n}} \right)^2 = \frac{1}{(2n)^2} \prod_{j=1}^{n-1} \left(1 + \frac{1}{2j} \right)^2 \geq \frac{1}{4n^2} n = \frac{1}{4n}.$$

Taking square roots on both sides gives

$$\binom{2n}{n} \frac{1}{2^{2n}} \geq \frac{1}{2\sqrt{n}}$$

which is the desired result.

²³Clearly, it holds for $n = 2$. For higher n , we get $\prod_{j=1}^{n-1} \left(1 + \frac{1}{j} \right) = \left(1 + \frac{1}{n-1} \right) \prod_{j=1}^{n-2} \left(1 + \frac{1}{j} \right) = \left(1 + \frac{1}{n-1} \right) (n-1) = n$ where the second equality uses the induction hypothesis.