Government Surveillance and Incentives to Abuse Power*

Paul Laskowski *Benjamin Johnson* *Thomas Maillart* *John Chuang*
University of California, Berkeley
School of Information
102 South Hall
Berkeley, CA 94720
{paul,johnsonb,thomas.maillart,chuang}@ischool.berkeley.edu

Abstract

We develop a stylized model that describes government surveillance, its effect on rational government actions, and the implications for citizens. Our model treats government as a rational actor that chooses how much to abuse its power, to prevent being voted out or overthrown. Starting with only limited directional assumptions, we prove a robust implication that increased surveillance results in increased abuse of power. Subsequently, we give specific functional relationships between surveillance, abuse, and popularity, all motivated by an underlying citizen-based utility model. We compute citizens' welfare in equilibrium, and find that when the government is unpopular, surveillance absolutely decreases welfare. However, in cases where the government is popular, we show that surveillance can increase welfare by lowering the risk of undesired change.

1 Introduction

In recent months, a series of news reports have brought public attention to the collection of individual data by government agencies. Classified documents, including those disclosed by Edward Snowden, have revealed a far more extensive surveillance apparatus than previously acknowledged, and fueled a public debate about the proper role of intelligence agencies in society. Proponents maintain that surveillance programs are instrumental in preventing terrorist attacks. Opponents counter that spying erodes individual privacy and facilitates totalitarian states. The public remains divided in the United States, even as politicians have initiated reforms aimed at protecting individuals from monitoring.

Network researchers have a central role to play in this process, as the technologies we work with facilitate, and in some cases impede, a range of surveillance techniques. While many researchers and policy makers would like to find a balance between security and privacy, we currently have little guidance for setting actual policies. Threats from both terrorists and abusive government officials are difficult to measure directly, and both are often kept secret from the public, though for different reasons.

Underlying the debate lies a set of research questions that are rooted in individual incentives: How does surveillance technology influence decisions by governments to abuse their power? How does it affect the odds of governmental change or revolution? What level of data gathering maximizes the welfare of a country's citizens?

This study will investigate these questions using a stylized game-theoretic model of surveillance. Our model centers around a government that wants to maximize its chances of remaining in power against an

*This research was supported in part by the National Science Foundation through award CCF-0424422 (TRUST - Team for Research in Ubiquitous Secure Technology). T. Maillart acknowledges support from the Swiss National Science Foundation (Grant No. PA00P2 145368).

opposition. It can choose to abuse its power in some way in order to damage the opposition. However, we will assume that a sufficiently abusive government will itself be disliked by its citizens.

Certainly, there is nothing inherently abusive about the use of surveillance. Private firms collect personal data to provide tools such as real-time traffic monitoring and product recommendations. Moreover many uses of surveillance by governments for homeland defense and law enforcement are uncontroversial. There are, however, many abuses of power that depend on data gathering and informational advantage for their effectiveness.

This fact can be witnessed in abuses ranging from blackmail, to speech suppression, to targeted violence. If the opposition is a politician, those in power can gain an advantage by wiretapping phones or monitoring email communications to find damaging material. When the opposition is a vocal political minority, the government can facilitate speech silencing by actively censoring critical comments and blocking access to sources of political information. If the opposition is an armed insurrection, governments can conduct secret assassinations using data provided by spy satellites, or they may even enact violence against citizens deemed to be sympathizers. The consequences of these actions can extend far beyond immediate targets, as even occasional violence can have a chilling effect on dissident speech and activities.

Stepping back from these associated abuses, our goal is to understand and ultimately evaluate the rational effect of surveillance technologies in our society. Recognizing that our topic invokes strong opinions on both sides, we will attempt to approach it dispassionately, without assuming that surveillance is inherently good or bad. Our framework incorporates several features to help us avoid bias and ground our results in first principles.

First, our technique makes a distinction between surveillance technology and the abuses of power that this technology may facilitate. Rather than assuming that governments are inherently civil or inherently corrupt, we merely assume that the choice to abuse power is influenced by incentives. In the spirit of game theoretic models, we will compute the rational level of abuse that a government would choose to maximize its odds of staying in power. Certainly, well-meaning individuals can avoid corrupt behavior, even when this conflicts with their self-interest. Nevertheless, we will interpret the rational level of abuse as a guide, useful for explaining increases or decreases in corrupt behavior.

Second, our model does not assume that a specific government is good or bad for its citizens. Instead, we recognize that a wide range of possibilities exists. A civil government will provide a high level of welfare to its citizens, making a change of government welfare-decreasing. On the other hand, a highly corrupt government imposes costs on its citizens, making a change welfare-enhancing.

Third, while our model will show that surveillance technology increases incentives to abuse power, it does not follow that the optimal level of surveillance is zero. Instead, we will highlight conditions under which surveillance is beneficial to citizens. Specifically, we will see that surveillance can enhance welfare by preventing the overthrow of a civil government. It is only when a government considerably abuses its power that surveillance becomes unfavorable, as it increases the effectiveness of that abuse and helps the government stay in power.

The rest of the paper is organized as follows. We review related work in Section 2. In Section 3, we introduce our general model that relies on concavity assumptions to derive relational results. We supplement the model in Section 4, with specific functional forms that are motivated by a rational model of citizens' behavior. We discuss several implications of our results and conclude in Section 5.

2 Related work

To the best of our knowledge, our paper is the first to directly model the impact of surveillance technologies on government incentives to abuse power.

Our model of citizen preferences builds on the framework of horizontal differentiation models that orig-

inates with Hotelling [7]. This framework was extended by Downs [5] to the arena of political competition. Further extensive literature in political science models the process by which governments are overthrown; however to the best of our knowledge this literature does not address the effects of surveillance. In the most similar study we are aware of, Ginkel and Smith [6] consider factors that determine the probability of revolution in a repressive regime. In a similar vein, Lohmann [13] describes the potential overthrow of a government through an informational cascade model. Kuran [9] attempts to explain why revolutions often take the world by surprise with a game theoretic model of political change.

A number of constraints on government surveillance in practice give context to our study. On the technological side, surveillance capacity is enhanced through scientific advancements in areas such as dataset aggregation [1]; while it is limited by anonymization technologies like Tor [4]. An ongoing debate surrounds the use of these and other technologies, in contexts ranging from ethics [3, 14], to law [2, 10, 11, 12], to security-relevant effectiveness [15]. While our modeling effort treats surveillance as a free parameter, our assumptions are compatible with these practical restrictions; so the debate surrounding use of surveillance by the government may be usefully informed by our incentive-based economic perspective.

3 General Model

In this section, we begin by defining the terms and variables used in our model. We introduce and motivate simple directional constraints about how these variables interact with each other. We conclude the section with proofs of two directional implications.

3.1 Terms and Variables

Government surveillance is a primary concept of interest; and our modeling framework addresses it directly through a surveillance level $S \in [0, 1]$. This parameter tries to capture the extent to which the government uses surveillance technologies to advance its interests. The surveillance level may be influenced by a number of factors including available technology, applications to security, legal restrictions, or administrative policy. To offer a meaningful interpretation, we propose that parameter condition $S = 0$ means that the government does not conduct any surveillance, while $S = 1$ means that the government watches everything. At non-extreme levels $S \in (0, 1)$, the government conducts some amount of limited surveillance and this amount increases with S.

Our second primary concept is the abuse of power by the government, and we address this through a parameter $A \in [0, 1]$. This parameter may be interpreted as a measure of the government's use of power *in excess of* moral or ethical standards of conduct. As the use of power is a multi-dimensional construct, individuals may differ in assessing which uses of power are more or less abusive. Some citizens may be more sensitive toward political corruption or back room dealing; while others may be more sensitive to physical violence or the repression of speech. To usefully employ a linear scale for abuse, we must project government actions onto a single dimension in a way that captures the concerns of most citizens. While this is a challenging problem, we note that similar instruments have been successfully used by other social scientists – for example the corruption perceptions index published by Transparency International [8]. We will use $A = 0$ to represent a government that never takes abusive actions. In contrast, $A = 1$ represents a government that entirely neglects ethical standards of conduct.

A central component of government's capacity to retain power is its base of support among its citizens. Our model includes the parameter $V_G \in \mathbb{R}$ which measures the government's popularity. For governments that are not democratic, we may also construe V_G alternatively as any other form of political capital such as a level of support from the country's military. As V_G decreases, the government becomes very unpopular and ultimately more likely to lose power, while as V_G increases, the government achieves uniform support to stay

in power. We will assume that V_G varies smoothly with both government surveillance S and government abuse A; and that the popularity of the government increases with V_G.

All governments have opposition, whether well-organized or not; and this opposition is also a key component of our modeling framework. To keep things simple at this point, we consider a single parameter $V_O \in \mathbb{R}$ which measures the popularity of the government opposition. We assume that V_O varies smoothly with government surveillance S and government abuse A; and that the popularity of the government's opposition increases with V_O.

Finally, all governments face a risk of losing power which we measure with parameter $p \in [0, 1]$. This probability p is a central component of the model, as the government's main objective is to choose an abuse level A^* which minimizes it.

For reference, all terms introduced in our basic modeling framework are listed in Table 1.

Table 1: Parameters in the Basic Modeling Framework

$S \in [0, 1]$	Level of surveillance by the government
$A \in [0, 1]$	Level of abuse of power by the government
$V_G \in \mathbb{R}$	Popularity of the government
$V_O \in \mathbb{R}$	Popularity of the opposition
$p \in [0, 1]$	Probability of power change

3.2 Directional Assumptions

When introducing the terms in our model, we included only assumptions related to the basic definitions of the terms. Here we argue for additional assumptions constraining the ways in which these terms relate to each other. We provide intuitive explanations for each of these directional assumptions.

Our first assumptions concerns the effect of small amounts of abuse on the relative popularity of government and its opposition. When the government is at the point committing no abuse at all, then a small additional amount of abuse, targeted at the government's opponent, will harm the popularity of the opponent more than it harms the popularity of the government. Together with the smoothness assumptions on V_G and V_O relative to A, this assumption formally implies

$$\frac{\partial(V_G - V_O)}{\partial A}|_{A=0} > 0. \tag{1}$$

Our second assumption contrasts the first by reversing the effect of abuse on popularity separation when the current abuse level is high. Informally, when the government becomes completely violent and corrupt, continued abuse yields less benefit to the government than to the opposition. Together with our smoothness assumptions on V_G and V_O mentioned previously, this assumption implies

$$\frac{\partial(V_G - V_O)}{\partial A}|_{A=1} < 0. \tag{2}$$

Our third assumption gives a plausible heuristic for interpolating between the extremes in the first two assumptions. Consider, for example, the government action of running a set of propaganda attack ads against a political opponent. In line with our first assumption, the very first such ad certainly harms the opponent's

popularity more than the government's. However, with increased ad exposure, the popularity-separating effect of each individual attack ad is less and less. To parallel this example, our general assumption here is that from any starting abuse level, each additional unit of abuse is less effective at separating the popularity of the government from that of the opposition. Formally, this assumption implies

$$\forall A, \frac{\partial^2}{\partial A^2}[V_G - V_O] < 0. \tag{3}$$

Our fourth assumption relates surveillance to the effectiveness of government abuse. The premise here is that surveillance allows the government to gather information which increases the effectiveness of power abuses. For example, surveillance may allow direct covert targeting of specific individuals, or even something more benign such as knowledge of a particularly effective communication strategy for accomplishing an objective. Other things being equal, for the same level of abuse, the government will be more effective at changing popularity in ways that are favorable to it, if the government has greater surveillance capacity. Together with smoothness conditions, this assumption implies

$$\forall S, \forall A, \frac{\partial}{\partial S}\frac{\partial}{\partial A}[V_G - V_O] > 0. \tag{4}$$

Finally, our fifth assumption concerns the probability of power change p. We assume that p is a smooth decreasing function of $V_G - V_O$. At the extremes, this assumption says that a government with no opposition retains power and a government that absolutely nobody supports will lose power. The smoothness condition gives a reasonably mild additional directional constraint on the probability of losing power; and it implies that

$$\frac{\partial p}{\partial (V_G - V_O)} < 0. \tag{5}$$

3.3 Directional Implications

In this Subsection, we prove two results related to the government's rational choice of abuse level. These results are derived using only the directional assumptions from the previous Subsection.

The first implication is that for any fixed level of surveillance, there is a unique optimal level of power abuse that minimizes the probability of the government losing power. More formally,

Result 1. *There is a unique power abuse level A^* satisfying*

$$\frac{\partial p}{\partial A}|_{A=A^*} = 0 \tag{6}$$

and p is at global minimum when $A = A^*$.

Proof. Since p is a decreasing function of $V_G - V_O$, we will minimize p by maximizing $V_G - V_O$. But from Equations (1), (2), and (3), $V_G - V_O$ is a twice differentiable function of A that is increasing at $A = 0$, decreasing at $A = 1$ and concave down everywhere. Such a function must have a unique maximum $A^* \in (0, 1)$, which is then a global minimizing value for the induced function $p(A)$. □

The second implication is that the government's optimal abuse level increases with its surveillance level. Formally,

Result 2. *Let A^* be the optimal level of power abuse by the government. Then we always have*

$$\frac{\partial A^*}{\partial S} > 0 \tag{7}$$

Proof. The basic intuition is that for any aggression level, increased surveillance makes abuse more effective; so that the government becomes more rationally motivated to use it. More formally, consider the graph of the function $\frac{\partial(V_G - V_O)}{\partial A}$ which starts above zero at $A = 0$ and ends below zero at $A = 1$. The optimal aggression level A^* is the point at which this curve crosses the x-axis. Now a small increase in surveillance moves every point on this curve upwards. Hence it necessarily moves the intersection point at which the graph crosses the axis to the right. That is to say, $\frac{\partial A^*}{\partial S} > 0$ □

4 Citizen-Based Model

In this section we supplement our modeling framework by specifying a specific functional relationship between the general model's terms. Motivating these relationships is an underlying citizen-based model, in which individual citizens respond game-theoretically to government actions.

4.1 Additional Terms and Variables

The underlying consumer model has citizens as actors. Each citizen has a type $i \in [0, 1]$. The type indicates a natural inclination to support change. So, a citizen of type $i = 1$ has a very strong natural inclination to support the government's opposition; while a citizen of type $i = 0$ is naturally inclined to support the current government.

A citizen's type i also determines its utility $u_i \in \mathbb{R}$. Each citizen makes a choice of whether to support the government or not, in order to maximize this utility. The government also has a utility $u_G \in \mathbb{R}$ which thus far is only constrained by the condition that it be strictly decreasing in p.

The demand for change $D \in [0, 1]$ is defined to be the proportion of citizen types that gain a higher utility from the opposition gaining power, than from the government staying in power. Finally, $L_G \in [-1, 1]$ is a parameter used to fine-tune the baseline likability of the government.

Table 2: Parameters Introduced in the Citizen-Based Model

$i \in [0, 1]$	Type of citizen (higher i implies more support for change)
$u_i \in \mathbb{R}$	Utility of a citizen with type i
$u_G \in \mathbb{R}$	Utility of the government
$D \in [0, 1]$	Citizens' demand for change
$L_G \in [-1, 1]$	Parameter used to adjust the baseline likability of the government

4.2 Additional Assumptions

Our first assumption in the citizen-based model is that the types of citizens are distributed uniformly over $(0, 1)$.

$$i \sim U(0, 1) \tag{8}$$

Our second assumption gives a formula for the utility of a citizen of type i. This utility depends on the type i, the relative popularity of the government to the opposition, and the event of whether or not a power

change occurs. Formally,

$$u_i = \begin{cases} V_G + (1 - i) & \text{if the government retains power} \\ V_O + i & \text{if the government loses power to the opposition} \end{cases} \tag{9}$$

The third assumption gives a formula for the utility of the government. This utility in the model is determined solely by the citizens' demand for change. That is,

$$u_G = -D \tag{10}$$

Note: we will show that this utility function is indeed strictly decreasing in p as stipulated in our more general assumptions.

Our fourth assumption gives a formula for the baseline popularity of the government. This popularity is determined by a baseline likeability parameter, plus the level of aggression against proponents of power change. In adherence to the concavity assumption in the relationship between popularity difference and abuse level, we have the government's popularity decreasing quadratically in the abuse level, and linearly with the demand for change. Specifically,

$$V_G = L_G - A^2 D \tag{11}$$

Finally, our fifth assumption gives a formula for the baseline popularity of the opposition. The opposition's popularity is reduced by the level of abuse directed against it, and this effect is enhanced with an increased surveillance level. Specifically,

$$V_O = -AS \tag{12}$$

4.3 Citizen-Based Model Implications

Here we prove a number of results related to the citizen-based model. The first result gives a direct formula for determining the citizens' demand for change in terms of the government's power abuse and surveillance levels. The next result finds the optimal level of power abuse to maximize the government's utility. The third result gives a formula for the rate of change of abuse with respect to surveillance, confirming that abuse increases with more surveillance. The last three results in this Subsection addresses welfare as a function of government surveillance.

Our first result gives a formula for the citizens' demand for change in terms of aggression and surveillance levels.

Result 3. *The citizens' demand for change can be expressed as*

$$D = \frac{1 - L_G - AS}{2 - A^2} \tag{13}$$

Proof. Recall that the demand for change D is the total mass of citizens whose utility is higher when the government is out of power. We determine this amount as follows.

The type i^* of a user who is indifferent between change and the government retaining power satisfies the equation

$$V_G + (1 - i^*) = V_O + i^*$$

which simplifies to

$$i^* = \frac{1 + V_G - V_O}{2}. \tag{14}$$

Any citizen whose type i is higher than i^* prefers change. Since the types are distributed uniformly over $(0, 1)$, the proportion of citizens who demand change is exactly

$$\begin{aligned} D &= 1 - i^* \\ &= 1 - \frac{1 + V_G - V_O}{2} \\ &= \frac{1}{2} - \frac{V_G - V_O}{2}. \end{aligned}$$

Substituting the values of V_G and V_O from Equations (11) and (12), we now obtain

$$\begin{aligned} D &= \frac{1}{2} - \frac{V_G - V_O}{2} \\ &= \frac{1}{2} - \frac{L_G - A^2 D + AS}{2} \\ 2D &= 1 - L_G + A^2 D - AS \\ (2 - A^2)D &= 1 - L_G - AS \\ D &= \frac{1 - L_G - AS}{2 - A^2}. \end{aligned}$$

□

Our second result gives a formula for the government's optimal abuse level A^* only in terms of the likability parameter L_G and the surveillance level S.

Result 4. *The government's optimal power abuse level can be expressed as*

$$A^* = \frac{(1 - L_G) - \sqrt{(1 - L_G)^2 - 2S^2}}{S} \tag{15}$$

Proof. The government maximizes its utility u_G by choosing an aggression level A that minimizes citizens' demand for change D. We compute the government's optimal choice by setting the partial derivative of u_G with respect to A equal to zero and solving for A.

$$
\begin{aligned}
0 &= \frac{\partial u_G}{\partial A} \\
0 &= \frac{\partial D}{\partial A} \\
&= \frac{\partial}{\partial A}\left[\frac{1 - L_G - AS}{2 - A^2}\right] \\
&= \frac{-S(2 - A^2) - (1 - L_G - AS)(-2A)}{(2 - A^2)^2} \\
0 &= -S(2 - A^2) - (1 - L_G - AS)(-2A) \\
&= -2S + A^2 S + 2A - 2AL_G - 2A^2 S \\
&= -A^2 S + 2A(1 - L_G) - 2S \\
0 &= A^2 S - 2A(1 - L_G) + 2S \\
A &= \frac{2(1 - L_G) \pm \sqrt{4(1 - L_G)^2 - 8S^2}}{2S} \\
A^* &= \frac{2(1 - L_G) - \sqrt{4(1 - L_G)^2 - 8S^2}}{2S} \\
&= \frac{(1 - L_G) - \sqrt{(1 - L_G)^2 - 2S^2}}{S}
\end{aligned}
$$

□

Our third result shows exactly how the government's rational abuse level changes with respect to its surveillance level.

Result 5. *Under our model's assumptions, the rate of change in the government's optimal abuse level relative to its surveillance level can be expressed as*

$$
\frac{\partial A^*}{\partial S} = \frac{1 - L_G}{S^2}\left(-1 + \frac{1 - L_G}{\sqrt{(1 - L_G)^2 - 2S^2}}\right) \tag{16}
$$

Proof. To save space in the derivation, we introduce the variable substitution $x = 1 - L_G$. Note that x is a constant with respect to S. Starting from Equation (15), we have

$$
\begin{aligned}
A^* &= \frac{x - \sqrt{x^2 - 2S^2}}{S} \\
\frac{\partial A^*}{\partial S} &= \frac{\left(-\frac{1}{2}(x^2 - 2S^2)^{-1/2}(-4S)\right) S - (x - \sqrt{x^2 - 2S^2})}{S^2} \\
&= \frac{\frac{2S^2}{\sqrt{x^2 - 2S^2}} - \left(x - \sqrt{x^2 - 2S^2}\right)}{S^2} \\
&= \frac{1}{S^2\sqrt{x^2 - 2S^2}}\left(2S^2 - x\sqrt{x^2 - 2S^2} + (x^2 - 2S^2)\right) \\
&= \frac{x}{S^2\sqrt{x^2 - 2S^2}}\left(-\sqrt{x^2 - 2S^2} + x\right) \\
&= \frac{x}{S^2}\left(-1 + \frac{x}{\sqrt{x^2 - 2S^2}}\right)
\end{aligned}
$$

The result now follows from reversing the variable substitution. □

Our last set of results concern welfare. In this context, welfare is the expected utility of all citizens. We assume that the number of individuals in government is small compared to the number of citizens, and therefore we do not include the government's utility in the welfare calculation.

Result 6. *Under our model's assumptions, the welfare of citizens can be expressed as*

$$w = V_G + \frac{1}{2} - p(V_G - V_O) \tag{17}$$

Proof. We compute welfare as follows. Given any fixed set of parameters, the world is in one of two states – either the government is still in power, or the opposition has taken over through change. In the world in which the government has power, the average utility of all citizens is computed by integrating over $i \in [0, 1]$ the function $u_i = V_G + 1 - i$, giving

$$\int_0^1 (V_G + 1 - i)di = \left[(V_G + 1)i - \frac{i^2}{2}\right]_0^1 = V_G + \frac{1}{2}.$$

In the world in which the government loses power, the average utility of all consumers is computed by integrating over $i \in [0, 1]$ the function $u_i = V_O + i$, giving

$$\int_0^1 (V_O + i)di = \left[V_O i + \frac{i^2}{2}\right]_0^1 = V_O + \frac{1}{2}.$$

The first case happens with probability $(1 - p)$, while the second happens with probability p. Therefore the expected value of consumer utility is

$$\begin{aligned} w &= (1 - p)\left(V_G + \frac{1}{2}\right) + p\left(V_O + \frac{1}{2}\right) \\ &= V_G + \frac{1}{2} - p(V_G - V_O) \end{aligned}$$

□

An interesting question we may ask about this model involves the effect of surveillance on welfare. To give a partial answer to this question, we may compute the derivative of welfare with respect to surveillance.

Result 7. *If the government is less likable than its opposition, then increased surveillance necessarily decreases welfare.*

Proof. Staring from Equation (17), we have

$$\begin{aligned} w &= V_G + \frac{1}{2} - p(V_G - V_O) \\ \frac{\partial w}{\partial S} &= \frac{\partial V_G}{\partial S} - \frac{\partial p}{\partial S}(V_G - V_O) - p\frac{\partial (V_G - V_O)}{\partial S} \end{aligned}$$

Using only directional assumptions from the baseline modeling framework , the terms on the right hand side in the equation above can be signed as follows:

$$\frac{\partial w}{\partial S} = (-) - (-)(?) - (+)(+).$$

The only term not unconditionally signed by our assumptions is the term $(V_G - V_O)$. If this term is negative, then the entire expression for $\frac{\partial w}{\partial S}$ is also negative. Consequently, if the opposition is more popular than the government, then the welfare will always decrease along with the surveillance level. □

On the other hand,

Result 8. *If the government is more popular than the opposition, it is possible for surveillance to increase welfare.*

Proof. Loosely speaking, this possibility requires two conditions to hold. First, the current government must provide citizens with considerably more utility than the opposition. Second, surveillance must quickly increase the probability that the government stays in power. Moreover, the marginal increase in power stability must be greater than the marginal decrease in the government's raw popularity advantage. To prove the possibility claim, it suffices to give an example of parameters that illustrate the desired effect. This example is provided in Figure 3, at the end of the next subsection.

□

4.4 Citizen-Based Model Example

Here we work through a full example in the citizen-based model to see how the increased availability of surveillance affects a range of other factors, including government abuse, raw support for the government, raw support for the government's opposition, the popular demand for power change, the popular support for the government, and citizen welfare. We provide graphs to illustrate how the model treats each of these factors, and we narrate a plausible scenario to go along with the effects in various regions of the parameter space.

In the following plots, we take the baseline likeability of the government to be $L_G = 0.25$; and the probability of power change as a function of demand to be a logistic function centered at $1/3$: $p(D) = \frac{1}{1+e^{100(1/3-D)}}$. Further, we only consider the effects of surveillance for values in the interval $[0, 1/2]$, because it is difficult to support a consistent narrative at the boundary conditions in which the optimal government abuse level reaches 1, or when the demand for change falls to zero. Mitigating these problems requires a more robust set of dynamics for our citizen-based framework. Within the presented parameter regions, the narrative and implications are clear.

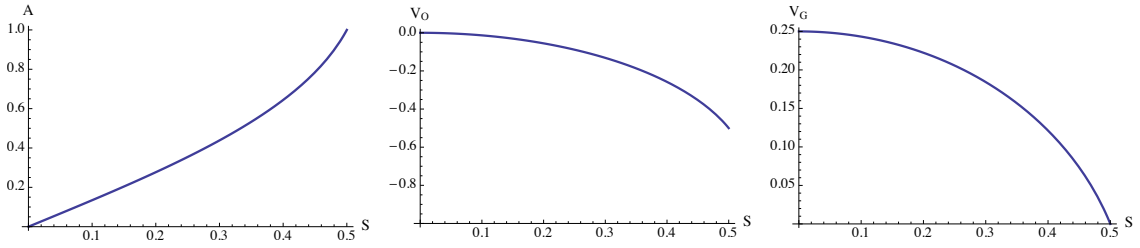

(a) Government Abuse Level (b) Opposition Support (c) Government Support

Figure 1: Effects of surveillance on government abuse and raw support levels with parameter settings: $L_G = 0.25$ and $p(D) = \frac{1}{1+e^{100(\frac{1}{3}-D)}}$

.

In Figure 1a, we see that the strategic government abuse level increases consistently with surveillance until it reaches its maximum value of 1. Figure 1b shows that support for the opposition falls steadily with increased surveillance. Figure 1c shows that support for the government also decreases with the surveillance level.

In Figure 2a, we see that the demand for regime change falls steadily with increased surveillance. Figure 2b shows the analogous effect in terms of the percentage of citizens who support the government. Fig-

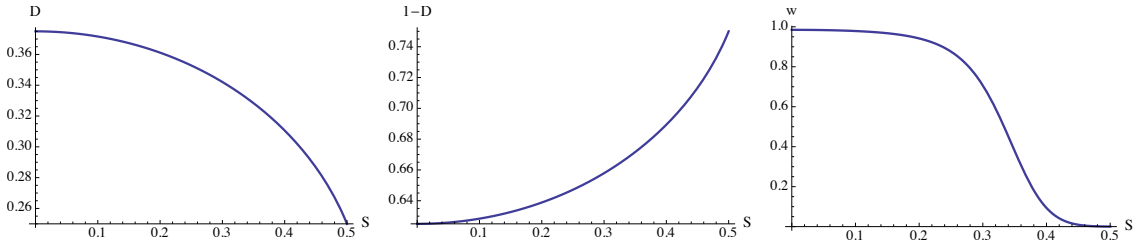

(a) Popular Support for Opposition (b) Popular Support for Government (c) Probability of Power Change

Figure 2: Effects of surveillance on citizens' demand for power change and the probability of change with parameter settings: $L_G = 0.25$ and $p(D) = \frac{1}{1+e^{100(\frac{1}{3}-D)}}$

ure 2c shows the probability of power change as a function of the surveillance level. This graph shows an abrupt drop in probability centered around a surveillance level of 0.35, which is when the demand for change is about $1/3$. This corresponds to the region of our logistic function p with the steepest slope. As the demand for change crosses that level, the probability of the government retaining power quickly approaches certainty.

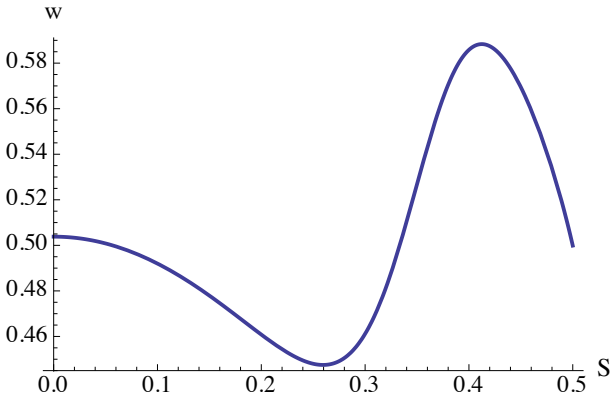

Figure 3: Welfare as a Function of Surveillance with parameter settings: $L_G = 0.25$ and $p(D) = \frac{1}{1+e^{100(\frac{1}{3}-D)}}$

Finally, Figure 3 shows how the model affects citizen welfare as a function of surveillance. We can use this graph to explain the effects of surveillance technologies on the model for an entire lifecycle, by considering its three directional regions.

As the surveillance level initially increases from zero, welfare decreases. In this region, the demand for change is well below 50%; yet the probability of change is high and only decreases slowly due to the flat region of the logistic function. Because the opposition is nearly certain to move into power, every citizen's utility is almost entirely determined by how much they support the opposition. Because increased

surveillance *decreases support for the opposition*, the utility of all citizens goes down, hence the decreasing welfare.

In the second region of the graph, welfare increases with surveillance. At this point, demand for change has become a bit lower, and the probability of change begins to decrease rapidly. The welfare increase can be entirely attributed to this *decreasing probability of regime change*. Most citizens prefer the government to stay in power, and because increased surveillance helps the government achieve this end, it also increases most citizens' utilities, hence the increase in welfare.

In the third region where welfare again decreases, the probability of power change has stabilized near zero, so that citizens are largely stuck with their current government. In this scenario, the dominant effect on citizen utilities comes from the *increasing abuse of power by the government*. This abuse drives down the popularity of the government, thereby decreasing almost every citizen's utility, hence the welfare decrease.

The narrative we provide is based on one choice of parameters satisfying our assumptions. Under alternative choices, some of the directional regions described above may not be present. Nevertheless, we believe that the competing effects we describe provide a useful lens for analyzing surveillance.

5 Discussion

"The means of defence against foreign danger have been always the instruments of tyranny at home." This worry was expressed by James Madison at the Federal Convention of 1787, but it is shared by many observers of today's government surveillance programs. In the intervening two hundred years, network technologies have transformed our society, but also enabled the collection of personal data on an unprecedented scale. While this ability has given us substantial benefits, including thwarted terrorist attacks and reduced crime, there is little doubt that it has also emboldened thuggish regimes around the world. As a new generation of technologies, including face-recognition software, highly-parallel databases, and natural language processing, stands poised to multiply the capabilities of those in power, the time is ripe to consider the effects of surveillance on how governments behave.

Though highly stylized, our model paints a more intricate picture of surveillance than suggested by either side of the public debate. Using nothing more than concavity assumptions, our most robust result shows that surveillance leads to greater abuses of power by government officials. At least in this context, it seems true that power corrupts. Once we added a rational model of citizen preferences, however, matters became less clear. When a government is unpopular, more surveillance absolutely decreases welfare. In cases where the government is well-liked, however, surveillance can increase welfare if it substantially lowers the risk of undesired change. Even so, a rational government will always want to increase surveillance, even beyond the level that's beneficial for citizens.

Our model is not without limitations. In an effort to maintain a reasonable level of simplicity, important concepts like abuse of power and surveillance were stripped of their complexity and reduced to single parameters. The government was represented as a single entity, and we did not include institutional checks on surveillance, such as the Foreign Intelligence Surveillance Court in the US. On the balance, we decided to abstract away from these details in order to focus on the basic relationship between surveillance and abuses of power.

Another limitation of our model was the assumption that citizen preferences are uniformly distributed between the government and the opposition. The uniform distribution is easy to work with but results in some unrealistic conclusions, like our prediction that a government can drive demand for change down to zero. Future extensions of this model might consider a more realistic distribution of citizens so that the demand for change is always positive. It may also be possible to calibrate the citizen model according to available statistics or new experimental measurements.

A third limitation of the model involves the treatment of the government's opposition as a single entity.

In fact, many governments risk losing power in multiple ways, for example, due to armed revolutionaries or by being voted out of office. Our model may be extended to understand surveillance towards different types of adversaries. Does surveillance against terrorists, for example, increase incentives for unethical behavior towards political opponents? In the same vein, we are planning to compare different types of surveillance: are population-wide dragnets more dangerous in terms of incentivizing corrupt behavior, than surgical data extraction?

As surveillance moves from the realm of spies and satellites to that of computer hardware, network scientists will find ourselves increasingly pivotal in mediating the balance of power between a state and its citizens. It is our hope that through modeling efforts like this one, we will gain tools to guide us in fulfilling this role's responsibility.

References

[1] Alessandro Acquisti and Ralph Gross. Predicting social security numbers from public data. *Proceedings of the National academy of sciences*, 106(27):10975–10980, 2009.

[2] Kevin Bankston and Ashkan Soltani. Tiny constables and the cost of surveillance: Making cents out of united states v. jones. *Yale Law Journal Online*, 123, 2014.

[3] Whitfield Diffie and Susan Landau. *Privacy on the Line: The Politics of Wiretapping and Encryption*. The MIT Press, updated and expanded edition edition, February 2010.

[4] Roger Dingledine and Paul F Syverson. *Privacy enhancing technologies*. Springer, 2003.

[5] Anthony Downs. An economic theory of democracy. 1957.

[6] John Ginkel and Alastair Smith. So you say you want a revolution a game theoretic explanation of revolution in repressive regimes. *Journal of Conflict Resolution*, 43(3):291–316, 1999.

[7] Harold Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929.

[8] Transparency International. Corruption perceptions index. `http://www.transparency.org/research/cpi/overview`.

[9] Timur Kuran. Sparks and prairie fires: A theory of unanticipated political revolution. *Public Choice*, 61(1):41–74, 1989.

[10] Susan Landau. National Security on the Line. *Social Science Research Network Working Paper Series*, April 2009.

[11] Susan Landau. Making sense from snowden: What's significant in the nsa surveillance revelations. *IEEE Security and Privacy*, 11(4):54–63, 2013.

[12] Susan Landau. Making sense of snowden, part ii : What's significant in the nsa revelations. *IEEE Security and Privacy*, 12(1):62–64, 2014.

[13] Susanne Lohmann. The dynamics of informational cascades. *World politics*, 47(1):42–101, 1994.

[14] David Lyon. *Surveillance as social sorting: privacy, risk, and digital discrimination*. Psychology Press, 2002.

[15] Bruce Schneier. It for oppression. *Security & Privacy, IEEE*, 11(2):96–96, 2013.