

# On the adoption of privacy-enhancing technologies

Tristan Caulfield<sup>1</sup>, Christos Ioannidis<sup>2</sup>, and David Pym<sup>1</sup>

<sup>1</sup> University College London  
t.caulfield@ucl.ac.uk, d.pym@ucl.ac.uk

<sup>2</sup> University of Bath  
c.ioannidis@bath.ac.uk

**Abstract** We propose a model, based on the work of Brock and Durlauf, which looks at how agents make choices between competing technologies, as a framework for exploring aspects of the economics of the adoption of privacy-enhancing technologies. In order to formulate a model of decision-making among choices of technologies by these agents, we consider the following: *context*, the setting in which and the purpose for which a given technology is used; *requirement*, the level of privacy that the technology must provide for an agent to be willing to use the technology in a given context; *belief*, an agent’s perception of the level of privacy provided by a given technology in a given context; and the *relative value* of privacy, how much an agent cares about privacy in this context and how willing an agent is to trade off privacy for other attributes. We introduce these concepts into the model, admitting heterogeneity among agents in order to capture variations in requirement, belief, and relative value in the population. We illustrate the model with two examples: the possible effects on the adoption of iOS devices being caused by the recent Apple–FBI case; and the recent revelations about the non-deletion of images on the adoption of Snapchat.

## 1 Introduction

Recent high-profile events — such as Snowden’s revelations about surveillance and the dispute between Apple and the FBI — have demonstrated the increasing significance of privacy concerns for individuals, organizations, and governments. As privacy-enhancing technologies become more widely available, and are increasingly incorporated into consumer products such as messaging apps, it is interesting and important to understand the factors affecting the adoption by consumers of different technologies. In this paper, we propose a model of how agents make choices between competing technologies, as a framework for exploring aspects of the economics of the adoption of privacy-enhancing technologies.

Acquisti et al. [3] deliver an excellent up-to-date survey of the economics of privacy. They provide for historical evolution of the economic theory of privacy from its early beginnings — starting with Posner [15,16] and Stigler [20], arguing in favour of limiting privacy in the name of market efficiency — to

the counterexamples where improved privacy (i.e., restrictions on the access to private information) may be welfare improving. According to Acquisti et al. [3]:

‘Privacy is, after all, a process of negotiation between public and private, a modulation of what a person wants to protect and what she wants to share at any given moment and in any given context.’

Other work has considered the role of privacy in technology adoption (for example, [17]) or considered economic factors affecting privacy [21] or privacy-enhancing technology adoption [1,2].

We introduce a characterization of privacy based on four key factors: *context*, the setting in which, and the purpose for which, a given technology is used; *requirement*, the level of privacy that the technology must provide for an agent to be willing to use the technology in a given context; *belief*, an agent’s perception of the level of privacy provided by a given technology in a given context; and the *relative value* of privacy, how much an agent cares about privacy in this context and how willing an agent is to trade off privacy for other attributes.

We introduce these concepts into the proposed model, admitting heterogeneity among agents in order to capture variations in requirement, belief, and relative value in the population.

In categorizing the agents’ different attitudes to privacy we adopt the useful classification of Harris and Westin [9,14], who divide the agents into three groups based upon their own perceptions of the value of their own privacy:

**The Fundamentalist.** Fundamentalists are generally distrustful of organizations that ask for their personal information, worried about the accuracy of computerized information and additional uses made of it, and are in favour of new laws and regulatory actions to spell out privacy rights and provide enforceable remedies. They generally choose privacy controls over consumer-service benefits when these compete with each other. About 25% of the public are privacy Fundamentalists.

**The Pragmatist.** Pragmatists weigh the benefits to themes of various consumer opportunities and services, protections of public safety or enforcement of personal morality against the degree of intrusiveness of personal information sought and the increase in government power involved. They look to see what practical procedures for accuracy, challenge and correction of errors the business organization or government agency follows when consumer or citizen evaluations are involved. They believe that business organizations or government should “earn” the public’s trust rather than assume automatically that they have it. And, where consumer matters are involved, they want the opportunity to decide whether to opt out of even non-evaluative uses of their personal information as in compilations of mailing lists. About 57% of the public fall into this category.

**The Unconcerned.** The Unconcerned are generally trustful of organizations collecting their personal information, comfortable with existing

organizational procedures and uses are ready to forego privacy claims to secure consumer-service benefits or public-order values and not in favour of the enactment of new privacy laws or regulations. About 18% of public fall into this category.

Sharing personal information may be perceived as risky or costly— facilitating identity theft, inviting unwanted attention by individuals or institutions, and possibly introducing limited participation in certain activities (e.g., exclusion from health insurance). Such negative impacts are known and the degree of aversion to the loss of privacy will differ between individuals depending upon their preferences and context.

The model — taking into account the privacy characteristics of competing technologies and the preferences of agents — indicates the expected levels of adoption of the competing technologies in different contexts. For example, sending different types of content with different levels of sensitivity over a service, such as Snapchat. By varying the parameters of the model — reflecting the characteristics of the technologies and the attitudes of the decision-making consumers — we explore how these factors influence the adoption of the different technologies.

In Section 2, we introduce the basic Brock-Durlauf model upon which our work is based. We also explain briefly our extension, from previous work [7], of this model to encompass multiple attributes. In Section 3, we present our main theoretical contribution. Using our analysis of the key characteristics of privacy, together with Westin’s characterization of attitudes towards privacy, we adapt our extended Brock-Durlauf set-up to model the adoption of privacy-enhancing technologies. In Section 4, we discuss two examples. First, the recent dispute between Apple and the FBI [4] and, second, Snapchat, exploring the effects of the population’s changing beliefs about and requirements for privacy. Finally, in Section 5, we summarize our analysis.

## 2 Background: the Brock–Durlauf model

Brock and Durlauf model a market where various technologies compete for adoption by a number of agents. The agents choose which technology to adopt based on the technologies’ relative profitabilities as well as the strength of the technologies’ social externalities; that is, how much the value of a technology increases as the number of other agents choosing it increases. This last feature makes the model particularly useful for looking at communication technologies— which form a large part of PETS—because the value of a technology increases with the number of people you can communicate with using it. The model can also look at exogenously-imposed policy, in the form of incentives or taxation, as well as increasing profitabilities through technological progress.

### 2.1 The basic Brock–Durlauf model

The basic model consists of  $M$  different technologies competing in a market for adoption by  $N$  agents. The utility for an agent of a technology  $\gamma$  in time period

$t$  is given by

$$u_{\gamma,t} = \lambda_{\gamma} + \rho_{\gamma}x_{\gamma,t} \quad (1)$$

where  $\lambda_{\gamma}$  is the profitability of the technology,  $x_{\gamma,t}$  is the fraction of agents using technology  $\gamma$  at time  $t$ , and  $\rho_{\gamma} > 0$  gives the strength of the social externalities. A low value of  $\rho_{\gamma}$  means the utility of the technology will not increase much as adoption rises; a high value means that the social component,  $\rho_{\gamma}x_{\gamma,t}$ , can influence the utility of the technology significantly.

In the model, each agent  $i$  experiences their own utility from their choice,  $\tilde{u}_{\gamma,i,t} = u_{\gamma,t} + \epsilon_{\gamma,i,t}$ , plus noise, where the noise term  $\epsilon_{\gamma,i,t}$  represents a random private component of utility and is independent and identically distributed across agents and known to the agent when it makes its decision. If the noise follows a double exponential distribution, then, as the number of agents tends to infinity, the probability that an agent will adopt technology  $\gamma$  at time  $t$ —which is equivalent to that technology’s share of the market—converges to

$$x_{\gamma,t} = \frac{e^{\beta u_{\gamma,t-1}}}{\sum_{j=1}^M e^{\beta u_{j,t-1}}}. \quad (2)$$

See [7] for more explanation of this equation.

The parameter  $\beta$  is inversely proportional to the variance of the noise,  $\epsilon$ , and characterises the degree to which choices made by the agents are determined by the deterministic components of utility. As  $\beta \rightarrow 0$ , choices are totally random and each technology will tend towards an equal share of the market; as  $\beta \rightarrow \infty$ , choices have no random component and the agents will all choose to adopt the technology providing the highest utility.

In Brock and Durlauf [5,6], the agents make decisions based on their expectations of the decisions of others in the same time period. The model can then be used to find the adoption equilibria. In contrast, we wish to look at the dynamics of adoption over time. Instead of using agents’ expectations about others’ decisions in the same time period, agents use information about the levels of adoption in the previous time period, as shown by the use of  $u_{c,t-1}$  in Equation 2.

The original definition of utility for a technology, in Equation 1, can be expanded to include a component determined by a policy-maker. This can represent, for example, some form of taxation or incentive designed to increase the adoption of a particular technology.

$$u_{\gamma,t} = \lambda_{\gamma} + \rho_{\gamma}x_{\gamma,t} - \tau_{\gamma}(x_{1,t}, \dots, x_{M,t}) \quad (3)$$

This policy component takes the form of a function,  $\tau_{\gamma}(x_{1,t}, \dots, x_{M,t})$ , for each different technology  $\gamma$  and gives the level of incentive or taxation based on the adoption shares of all the technologies in the market. This means that, for example, a policy-maker could apply an incentive to a technology that decreases as it becomes more widely adopted. Policies that tax one technology and use the benefits to promote another can be modelled by using opposite-signed functions on the two technologies.

Switching costs can also be added to the model by introducing asynchronous updating. That is, a portion  $\alpha$  of the agents do not switch technologies in each time period, simulating the retarding effect switching costs have on the speed with which new technologies are adopted:

$$x_{\gamma,t} = \alpha x_{\gamma,t-1} + (1 - \alpha) \frac{e^{\beta u_{\gamma,t-1}}}{\sum_{j=1}^M e^{\beta u_{j,t-1}}}. \quad (4)$$

**Equilibria.** The model allows for equilibria; that is, where the share of adoption in one time period is the same as the previous time period. For low values of  $\rho$ , there will only be one equilibrium point. For higher values, it is possible to have multiple equilibria. In general, the model will, over time, approach one of the equilibrium points.

Except in the case where  $\beta = \infty$ , a technology will never have all of the share of the market or become extinct: some (possibly very small) portion of the population will continue to choose it.

## 2.2 Extension to multiple attributes

In [7], we looked at how the Brock–Durlauf model could be applied to the adoption of encryption technologies. A key point from this work is that representing technologies with a single attribute, profitability, is not suitable for creating useful models about encryption adoption. Instead, it is necessary to use multiple attributes which better represent the technologies and the way decisions to use them are made. Multi-attribute utility theory is explained in [13] and applied to security in [11].

This is achieved by adapting the model to use a set of attributes,  $A$ . Now, the utility for each technology (Equation 1) becomes

$$u_{\gamma,t} = \sum_{a \in A} v_{\gamma,a} + \rho_{\gamma} x_{\gamma,t}, \quad (5)$$

where  $v_{\gamma,a}$  is the value of attribute  $a$  for technology  $\gamma$ .

Similarly, including policy, Equation 3 becomes

$$u_{\gamma,t} = \sum_{a \in A} v_{\gamma,a} + \rho_{\gamma} x_{\gamma,t} - \tau_{\gamma}(x_{1,t}, \dots, x_{M,t}). \quad (6)$$

The attributes used depend on the technologies being modelled and the purpose for which the models are intended. In [7], we used three attributes: monetary cost, functionality, and usability.

## 3 Modelling privacy

The basic approach of the model as described in Section 2 is not adequate for modelling the adoption of privacy-enhancing technologies. The model must be

extended to capture the characteristics of privacy. This represents a significant enrichment of the model to capture a more complex collection of interacting factors, including heterogeneity of agents.

In this section, we first discuss these characteristics; then, we describe how the model is extended to include them. Finally, we discuss the effects of different choices of parameters.

### 3.1 The characteristics of privacy

We consider a society of decision-making entities who wish to protect the privacy of certain information that they own in a range of contexts in which they interact with the providers of goods and services. These interactions are typically enabled by technologies with differing privacy-protecting characteristics.

Some transactions are more sensitive than others for some individuals. For example, some individuals will choose to use online banking services, in which private information is potentially exposed to the public internet, and some will prefer to perform their financial transactions in person at a branch of their bank, where the immediate exposure is limited to the specific bank employees involved.

We can deconstruct this situation in different ways. It may be the user of online banking simply does not place a high value on their privacy or it may be that they do place a high value on their privacy, but also believe that the bank's systems provide adequate protection for their judgement of value of their privacy. Similarly, the in-branch may believe that the online privacy protections do provide adequate protection for their judgement of the value of their privacy.

This set-up illustrates two characteristics that we need to incorporate into our model: first, that agents have a judgement of the value of their privacy; and, second, that they have beliefs about the ability of a given technology to protect their privacy given their judgement of its value.

These two examples illustrate the use of particular technologies to access services in specific contexts. In general, services, such as banking, will be accessed in different contexts. For example, the user of online banking may be willing to use the service from a personal computer at home, but not from a shared public computer: their belief about the level of protection is dependent on the context.

So, in order to formulate a model of decision-making among choices of technologies by these agents, we must consider what are the relevant characteristics of privacy in this context.

- *Context*: the setting in which and the purpose for which a given technology is used.
- *Requirement*: the level of privacy that the technology must provide for an agent to be willing to use the technology in a given context.
- *Belief*: an agent's perception of the level of privacy provided by a given technology in a given context.
- *Relative value of privacy*: how much an agent cares about privacy in this context and how willing an agent is to trade off privacy for other attributes

Attitudes to privacy have been classified into three groups — fundamentalist, pragmatist, and unconcerned — by Westin [9,22]. The final characteristic above, the relative value of privacy, includes the idea of a trade-off between privacy and other attributes. The Westin groups provide a convenient way in which to organize agents into groups with similar trade-off preferences. The examples in Section 4 illustrate this organization.

### 3.2 An adoption model using the privacy characteristics

We can capture these characteristics of privacy in the model by making some changes to its structure.

First, we can capture context by increasing the granularity of the model— instead of looking at technologies’ share of the market, we can look at how adoption is shared between technologies’ use in different contexts. Each technology is divided into multiple technology–context components, and the model now looks at how agents choose between these.

We introduce a set of all of these components,  $C$ , with subsets  $C_\gamma$  containing all of the components for technology  $\gamma$ . Now, we define  $u_{c,t}$  to be the utility of *component*  $c$ , rather than a *technology*. Similarly,  $x_{c,t}$  is now the share of a component, not a technology, at time  $t$ .

The total share of a technology  $\gamma$  is now given by the sum of its components:

$$x_{\gamma,t} = \sum_{c \in C_\gamma} x_{c,t}. \quad (7)$$

As an example, consider a cloud storage technology, where users can keep backups of their files. This could be divided into three different contexts based on its use for different purposes: storing photos, storing documents, and using it to do both of these. Each context offers different advantages (and so has different values for its attributes), and for each context agents may have different requirements for privacy. One agent might feel that photos require more privacy than documents do, whereas another might feel the opposite.

In the model up to this point, agents have been homogenous in terms of the utility they receive from a technology, with the only difference coming from  $\epsilon_{c,i,t}$ , the private utility they receive from the noise. Modelling privacy requires heterogeneity: each agent has different preferences towards privacy, different requirements, and a different willingness to trade privacy for other attributes.

We add this to the model by giving each agent  $i$  a value  $b_{c,i} \in [0, 1]$  for their belief about how well a component preserves or provides privacy, a value  $r_{c,i} \in [0, 1]$  for the agent’s required level of privacy for a component, and a value  $w_{c,i} > 0$  as a weight indicating the relative importance of privacy (in a component) to other attributes in the model. The utility function used in the model then becomes

$$u_{c,i,t} = \pi_{g(i)}(b_{c,i}, r_{c,i}, w_{c,i}) + \sum_{a \in A} v_{c,a} + \rho_c x_{c,t} - \tau_c(x_{1,t}, \dots, x_{M,t}). \quad (8)$$

where  $g(i)$  gives the group an agent belongs to and  $\pi_g(i)(b_{c,i}, r_{c,i}, w_{c,i})$  is a tradeoff function that specifies how the utility an agent receives from privacy changes for varying levels of belief, requirement, and value of privacy—essentially, how willing they are to trade privacy for other attributes.


Introducing the idea of a group here provides a convenient way of representing different attitudes towards security, and allows us to capture ideas such as Westin’s [22,9] groups. In theory, each agent could belong to its own group, each with a different trade-off function, but it would be immensely difficult to get the data required to fit a function to each participant in a study, for example. Agents in a group share the same trade-off function, meaning that they respond to different values of belief and requirements about privacy in the same way.

In this paper, we divide the population of agents into three groups, based on Westin’s classifications of attitudes about privacy. Each group has a different trade-off function, which are shown in Figure 1. For those unconcerned about privacy, there is little difference between components that meet requirements and those that do not. For pragmatists, any component that satisfies requirements receives the full utility value, with a linear trade-off for those that do not. For fundamentalists, there is very steep decline in utility value—quickly going negative—for components for which beliefs about privacy do not meet requirements. The trade-off functions are

$$\pi_{fund}(b_{c,i}, r_{c,i}, w_{c,i}) = w_{c,i} \frac{0.5 + \tanh(10(b_{c,i} - r_{c,i} + 0.1))}{1.5} \quad (9)$$

$$\pi_{prag}(b_{c,i}, r_{c,i}, w_{c,i}) = \begin{cases} w_{c,i} & b_{c,i} - r_{c,i} > 0 \\ w_{c,i}(b_{c,i} - r_{c,i} + 1) & b_{c,i} - r_{c,i} \leq 0 \end{cases} \quad (10)$$

$$\pi_{unco}(b_{c,i}, r_{c,i}, w_{c,i}) = 0.1w_{c,i}(b_{c,i} - r_{c,i}) + 0.9. \quad (11)$$



**Figure 1.** Trade-off functions for each of the Westin groups. The figure shows the utility received from privacy given the difference in beliefs and requirements when privacy value is 1.



We can update Equation 3 to account for the heterogeneity by summing over the population of agents. Now, the share of each technology–context component is given by


$$x_{c,t} = \frac{1}{N} \sum_{i=1}^N \frac{e^{\beta u_{c,i,t-1}}}{\sum_{j=1}^C e^{\beta u_{j,i,t-1}}}. \quad (12)$$

Each agent here has an equal weight ( $1/N$ ) and represents an equal share of the population, but this could easily be changed so that agents have different weights, making them representative of different proportions of the population. This might be useful, for example, when polling a population, where some agents’ characteristics have a greater likelihood.


### 3.3 Parameters

**Sample size.** In the examples below, we approximate the distribution of preferences about privacy, including requirements, beliefs, and values by using beta distributions to represent the distribution of values in the population. We then sample from these distributions to create a collection of agents with heterogeneous properties.


As the sampling is random, the points chosen can influence the behaviour of the model. We ran 100 trials for each of a number of different sample sizes in order to observe the magnitude of this influence. Figure 2 shows the mean and  $\pm 2\sigma$  values for each of the sample sizes.



**Figure 2.** Demonstration of the effect of sample size. The figure shows the mean and  $\pm 2\sigma$  for samples sizes of 100, 500, 1000, 2500, 5000, and 10000.



**Figure 3.** Demonstration of the effect of parameter  $\beta$ . With a high  $\beta$ , the adoption of the more profitable technology is greater.



**Figure 4.** Demonstration of the effect of parameter  $\rho$ .

As expected, the variance of the low sample sizes is higher than for the larger sample sizes. For 100 samples, the 5%-95% range is 0.03491; for 5000 it is 0.0053, and for 10000 samples it is 0.0030. We use 10000 samples in all examples below.

**The  $\beta$  parameter.** The parameter  $\beta$  is inversely related to the variance of the noise  $\epsilon_{\gamma,i,t}$ , which is a private component of the utility an agent gets for a particular choice of technology  $\gamma$ . As the variance of the noise grows—so, as  $\beta$  grows smaller—the less the other, deterministic components of utility matter. Conversely, as  $\beta$  grows larger and the variance of the noise decreases, agents increasingly make their choice based on the deterministic parts of the utility function.

Figure 3 shows the adoption over time for a technology for different values of  $\beta$ . The technology shown is slightly more profitable than the competing technology, but all other values are the same. For low  $\beta$ , the more profitable technology shares the market with its competitor evenly. As  $\beta$  grows, more agents adopt the more profitable technology.

In the examples below, we use a value  $\beta = 3.0$ .

**Social effects.** The parameter  $\rho_c$  controls the strength of social effects for a component. Figure 4 shows the effect of different values of  $\rho$  on the adoption curve (plotting  $x_t$  against  $x_{t+1}$ ) of a technology. Both technologies use the same parameter values (including  $\rho$ ), except for profitability, which is slightly higher for the technology shown.

As the value of  $\rho$  grows, the utility of a technology increases with increased adoption. High values of  $\rho$  amplify increases in utility from adoption.

Also shown in the figure is a diagonal line where  $x_t = x_{t+1}$ , meaning that the system is in equilibrium. For the lower values of  $\rho$  the adoption curves only have one equilibrium, meaning that adoption will approach this point over time. When  $\rho = 1$ , there are three equilibria: two stable points, low and high, and

a third, unstable point near  $x = 0.4$ . If the initial state is below the unstable equilibrium, adoption will move towards the lower equilibrium; if it is higher than the unstable equilibrium, adoption will move towards to the higher stable equilibrium.

## 4 Examples

We discuss in detail two examples. First, we consider the recent dispute between Apple and the FBI [4], with the purpose of demonstrating how beliefs and requirements about privacy influence adoption. Second, we consider Snapchat ([www.snapchat.com](http://www.snapchat.com), accessed 03/03/2016), a picture-messaging app which promised that images were available only for brief periods time, but for which it transpired that images were in fact retained [18]. We use this example to demonstrate the role of context in privacy decision-making regarding the use of the technology.

Both of these examples are intended to be illustrative of the theoretical model. It would of course be valuable to condition the examples on empirical data. However, such data collection and analysis, requiring substantial dedicated effort, is beyond our current scope. The paper by Brock and Durlauf [6] shows how the basic model can be fitted by maximum likelihood estimation; in principle, our extensions can be given a similar analysis.

Building on our discussion in Section 3, we remark that the example discussed there — namely access to banking services — would also provide an examples of the issues discussed in this section.

The model is implemented using the julia language [12].

### 4.1 Apple v FBI


In California, there is an ongoing case between Apple and the FBI, where the FBI is investigating the San Bernardino killings and wishes to access one of the killer’s locked and encrypted iPhones. The FBI is seeking a court order, under the All Writs Act, to compel Apple to assist in unlocking the device, possibly by creating and signing a custom firmware image that would allow the FBI to brute-force the password easily. Apple has argued against the FBI and the case has generated a great amount of media coverage.

For this example, we are interested in the effects this media coverage. Apple has publicly stated during the course of the case that it believes firmly in the privacy of its customers; this can be viewed as a strong signal about the level of privacy provided by Apple products and agents may update their beliefs in response, resulting in a change of technology choice. We will use the model to explore how adoption changes in response to shifting beliefs about the privacy a technology provides, shifting requirements, and shifts in both of these simultaneously.


**Set-up.** In this example, we look at two technologies competing against each other, without considering any distinct contexts. The first technology is Apple’s iPhone, and we look at its adoption when competing against Android phones.

For simplicity, we do not consider any attributes other than cost in this example; we assume that usability and functionality are essentially equivalent between the devices. Cost, on the other hand, differs: Apple devices tend to be more expensive than the bulk of android phones. Accordingly, we use a value of 1.1 for Apple and 1.5 for Android.

The value of  $\rho$  indicates how much the utility agents gain from adopting a technology increases as more agents begin to use it. In the case of mobile phones, they are largely interoperable with each other, and many of the applications written for them are present on both Apple and Android devices, suggesting that the value of  $\rho$  should be low. However, there are functions on the phone, such as Apple’s iMessage, which increase in utility as more people use them, meaning that there is some social effect present. For this example, then, we use a value of  $\rho = 0.5$  for both technologies.




**Figure 5.** Distributions representing the value of privacy for the different Westin groups in the population.




**Figure 6.** Distributions representing the requirements for privacy of the different Westin groups in the population.

We need to make some assumptions about the distributions of values, beliefs, and requirements of security in the population. First, Figure 5 shows the distributions we are using for the value of privacy. There is a separate distribution for each of the three Westin categories. We assume that fundamentalists are more likely to place a higher value on privacy than the pragmatic and the unconcerned. Similarly, we assume that it is more likely that the pragmatic have a higher value of privacy than the unconcerned. For this example, we say that privacy fundamentalists form 25% of the population, pragmatists 55%, and the unconcerned the remaining 20%.

Next, Figure 6 shows the distributions from which requirements about privacy are drawn. Again, we assume that fundamentalists are likely to have higher requirements than the pragmatic, and the pragmatic are likely to have higher requirements than the unconcerned.



**Figure 7.** Distributions representing the initial beliefs about the privacy provided by Apple and Android mobile phones.



**Figure 8.** Increasing beliefs about the privacy provided by Apple phones.

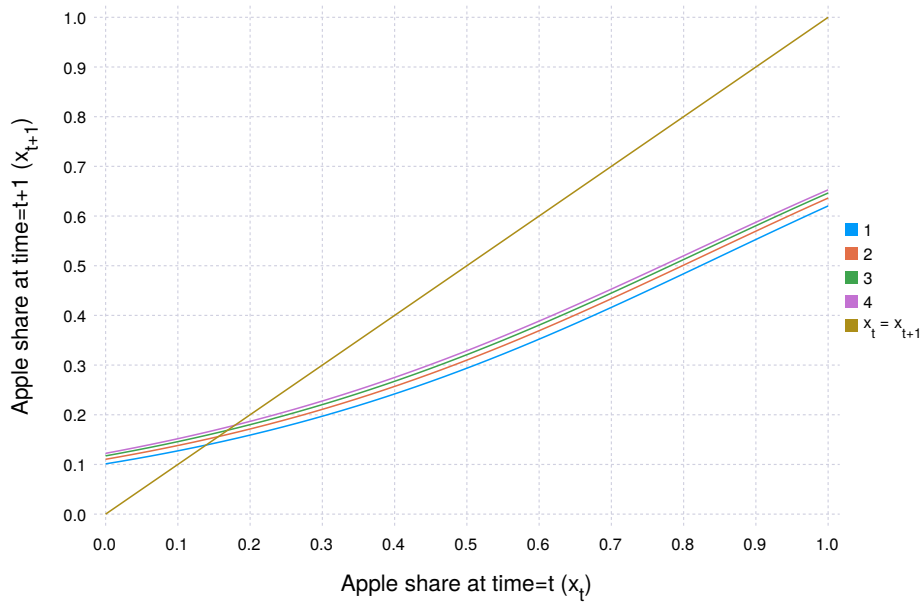
Finally, we look at the distributions from which we sample values for belief about the privacy provided by the different technologies. These distributions of belief are shared by the entire population and are not segmented into Westin groups. Figure 7 shows the distributions; agents are more likely to believe that an Apple phone provides a greater level of privacy compared to Android than vice versa.

**Changing Beliefs.** Now, we will examine what are the likely effects on adoption of a shift in beliefs about the privacy provided by Apple phones. As stated above, the shift is a hypothetical one caused by the media attention around the Apple v FBI case and Apple’s public stance on privacy. As such, we will look at how adoption changes for different magnitudes in shifts in belief to understand the range of possible effects.

We model the shifts in beliefs by changing the distribution of beliefs in the population and randomly sampling again. We look at four different distributions of beliefs about the privacy of Apple phones; we do not alter the distribution for Android phones. The different distributions are show in Figure 8, labeled 1–4, each with increasing probability of a higher belief about privacy. The first is the same distribution shown in Figure 7.

**Table 1.** Equilibrium Apple share values for shifts in belief.

Shift	Equilibrium
1 (orig.)	0.1402
2	0.1585
3	0.1667
4	0.1806



**Figure 9.** Adoption curves for the different levels of belief.


The resulting adoption curves are shown in Figure 9. The shifts in belief about the privacy provided by Apple phones result in increased adoption. Table 1 shows the equilibrium values for the four shifts. The base case, 1, shows Apple with a 14% share of the market—intentionally close to the actual share in 2015 [10].

With each shift, the share of the market grows, showing that agents receive greater utility from technology that better meets their requirements and thus switch.

**Changing Requirements.** Next, we consider what happens if the media coverage increases agents’ awareness of the need for privacy, resulting in a shift in requirements. As we are using different distributions of requirements for each of the Westin categories, we need to shift all three distributions to model the change in requirements. These are shown in Figure 10. In each case, the distributions shift to the right, making higher values for the requirement for privacy more likely.

The adoption curves for the shifts in requirement are shown in Figure 11. Unlike the shifts in beliefs, the shifts in requirements do not result in increased adoption. As Table 2 shows, there is a fractional increase, indicating that some agents are switching technologies, but this could also be explained by sampling variance.

This behaviour is expected, when considering the way the model is constructed. The previous shift in belief change the value for just Apple technology,



**Figure 10.** Shifting requirements for each of the Westin groups.


**Table 2.** Equilibrium Apple share values for shifts in requirement.

Shift	Equilibrium
1 (orig.)	0.1417
2	0.1429
3	0.1451
4	0.1429

increasing its utility. This shift in requirements changes the requirements for both Apple and Android, meaning that any relative changes will be smaller. Agents of the pragmatic or unconcerned types will not experience a large relative change in utility when requirements shift—any increase for Apple is likely to be too small to overcome the utility derived from cost. The only fundamentalist agents that would change technologies are those for whom both technologies met their requirements before the shift and only Apple after the shift.

**Changing Beliefs and Requirements.** Here, we look at what happens if there are shifts in both belief and requirement simultaneously. We use the same shifts as previously shown in Figure 8 and Figure 10.

Figure 12 shows the adoption curves when both belief and requirements are shifted. The equilibrium values are shown in Table 3. The increase in adoption here is greater than in the shift of beliefs or requirements alone. The combination



**Figure 11.** Adoption curves for the different levels of requirements.

**Table 3.** Equilibrium Apple share values for shifts in belief and requirement.

Shift	Equilibrium
1 (orig.)	0.1400
2	0.1634
3	0.1840
4	0.2011


of shifting both beliefs and requirements results in a relative increase in utility for Apple.

## 4.2 Snapchat


In this example, we explore the use of contexts and how privacy affects in which contexts agents choose to use technology by looking at the ephemeral picture messaging application Snapchat. This is a widely used application that gives its users the ability to control how long the messages they send can be seen by the recipients, after which the messages are deleted. However, the messages are not deleted securely, and can still be recovered after they have disappeared from the application.

**Set-up.** Roesner et al. [18] survey users of Snapchat, asking which types of content users send and how they feel about privacy. They give the breakdown of the





**Figure 12.** Adoption curves for increasing levels of belief and requirements.



**Figure 13.** Adoption over time for different contexts. The revelation about Snapchat’s non-deletion of messages occurs after time  $t = 5$ .

study participants into Westin groups (39.4% fundamentalist, 45.7% pragmatist, 12.6% unconcerned) and report how many users *primarily* send sexual content (1.6%) and how many *have sent* sexual content (14.2%).

We use these values directly in this example. We model three different contexts: sending only explicit content, sending only non-sensitive content, and using the technology for both. We say that Snapchat is competing against another similar application (which has the same contexts), but Snapchat initially has the majority share of adoption, around 90%.

We assume that the values of usability and cost are the same for both technologies, but there is a difference in the utility received from functionality. For Snapchat, we assign using it for only explicit content the value 0.9; for mixed explicit and non-sensitive use 1.5, and for non-sensitive use only 1.54. For the competing technology, in the same order, we use the values 0.8, 1.2, and 1.44. These values were chosen so that the model roughly matches the values reported in Roesner et al. [18]. The values for the explicit-only and mixed-content use contexts are less than the non-sensitive context. This is because—even though an agent using the technology for both types of content technically has greater functionality—the proportion of agents who actually generate explicit content is very small and the attribute values reflect the utility received by the *population* of agents.

Since we are looking at messaging applications, the value of social effects is very high: the utility of such an application increases with the number of people you can contact using it. As such, we use a value of  $\rho = 1$ .

The distributions used for beliefs, requirements, and values are the same initially for the two technologies. Fundamentalists are likely to have very high requirements and to place a high value on privacy for the explicit and mixed-content messaging contexts, and higher-than-average requirements for non-sensitive messaging. The unconcerned have the lowest requirements and values, and the pragmatists are in between the two other groups.

**Change in beliefs.** We model the revelation that Snapchat messages are not securely deleted and can be recovered as a shock which causes a downward shift in belief about the privacy provided by Snapchat. Beliefs about the competing product do not change.

Figure 13 shows the share of adoption of the various components of Snapchat and its competitor over time, as well as the total market share for Snapchat. The shock occurs after time  $t = 5$ .

**Table 4.** Adoption of different components before and after Snapchat’s non-deletion of messages is revealed.

	Before	After
Snapchat explicit	0.014	0.011
Snapchat mixed	0.119	0.094
Snapchat non-sensitive	0.742	0.764
Comp. explicit	0.014	0.028
Comp. mixed	0.038	0.032
Comp. non-sensitive	0.074	0.071
Total Snapchat	0.874	0.869

The initial values, before the shock, are close to the values reported in Roesner [18]. Out of Snapchat’s share—not the total share—1.5% use it for explicit messages only, compared to 1.6% in Roesner, and 13.2% use it for mixed content, compared to 14.2%.

Table 4 shows the values of adoption for the different components before and after the shock. The use of Snapchat for explicit messaging decreases from 1.4% to 1.1%. Similarly, the use of Snapchat for mixed explicit and non-sensitive messaging declines from 11.9% to 9.4%. The use of Snapchat in a non-sensitive context actually increases, from 74.2% to 76.4%, showing that agents who have a high level of privacy requirement in the explicit or mixed-content messaging contexts no longer use the technology in those contexts when they believe that their privacy requirements are no longer being met.

Snapchat’s total share declines post-shock from 87.4% to 86.9%. The agents that switched technologies to the competitor did so for explicit messaging, which grew from 1.4% to 2.8%. The beliefs about the security of the competing product did not change, so agents wishing to use the technology for explicit content were willing to switch to a product with less functionality that met their privacy requirements.

## 5 Conclusions

We have discussed the characteristics of privacy from the point of view the economic agent:

- *Context*: the setting in which and the purpose for which a given technology is used;
- *Requirement*: the level of privacy that the technology must provide for an agent to be willing to use the technology in a given context;
- *Belief*: an agent’s perception of the level of privacy provided by a given technology in a given context;
- *Relative value of privacy*: how much an agent cares about privacy in this context and how willing an agent is to trade off privacy for other attributes.

We have incorporated these characteristics into a model of technology adoption by a society of heterogenous decision-making agents.

Our analysis is based on Harris and Westin’s classification of agents as Fundamentalist, Pragmatist, and Unconcerned. For each of these groups, we have assigned a function that determines the utility an agent derives from a technology, depending upon the agent’s beliefs about how effectively the technology meets their requirements for protecting their privacy.

We have presented two main examples. First, to demonstrate the effects of changing beliefs and requirements, we have considered the signal of concern for privacy suggested by ongoing Apple v FBI dispute. Second, we have demonstrated the model’s use to capture context by considering the change in types of messages that are prevalent on Snapchat before and after a change in beliefs about the level of privacy provided.

The literature on economic modelling of privacy and its role in technology adoption is quite limited, with [8] and the references therein providing a good guide. We believe the present paper represents a useful contribution in that we identify key characteristics, create a model that is capable of capturing them, and explore, with examples, their significance.

The model we have presented here allows preferred attributes for particular agents to be specified. Future work might employ empirical studies of the preferences, beliefs, and requirements of actual agents and incorporate this data into the model. Similarly, the trade-off functions used for the Westin groups might be derived from empirical studies.

The model as presented includes a policy component that is not exploited in this paper. Further work might explore the role of policy in the adoption of privacy-enhancing technologies.

## References

1. A. Acquisti. Protecting privacy with economics: Economic incentives for preventive technologies in ubiquitous computing environments. *Proc. of Workshop on Socially-informed Design of Privacy-enhancing Solutions*, 4th UBICOMP, 2002.
2. A. Acquisti and J. Grossklags. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy* 3(1), 2005, 26–33.
3. A. Acquisti, C. Taylor, and L. Wagman. The Economics of Privacy. <http://ssrn.com/abstract=2580411> (visited 02/03/2016).
4. FBI–Apple encryption dispute. [https://en.wikipedia.org/wiki/FBI-Apple\\_encryption\\_dispute](https://en.wikipedia.org/wiki/FBI-Apple_encryption_dispute) (visited 13/5/2016).

5. W.A. Brock and S.N. Durlauf. Discrete Choice with Social Interactions. *The Review of Economic Studies*, 68(2):235–260, 2001.
6. W.A. Brock and S.N. Durlauf. A Multinomial-Choice Model of Neighborhood Effects *Amer. Econ. Rev.* 92(2):298–303, 2002.
7. T. Caulfield, C. Ioannidis, and D. Pym. Discrete Choice, Social Interaction, and Policy in Encryption Technology Adoption [Short Paper] In *Proc. Financial Cryptography and Data Security '16*, J. Grossklags and B. Preneel (editors), Lecture Notes in Computer Science, forthcoming, 2016.
8. ENISA. Study on monetising privacy: An economic model for pricing personal information. 2012. <https://www.enisa.europa.eu/activities/identity-and-trust/library/deliverables/monetising-privacy> (visited 04/03/2016).
9. L. Harris and A. Westin. The Equifax Canada Report on Consumers and Privacy in the Information Age. Technical Report, 1992.
10. IDC: Smartphone OS Market Share 2015, 2014, 2013, and 2012. <http://www.idc.com/prodserv/smartphone-os-market-share.jsp> (visited on 04/03/2016).
11. C. Ioannidis and D. Pym and J. Williams. Investments and Trade-offs in the Economics of Information Security. In: *Proc. Financial Cryptography and Data Security '09* (R. Dingledine and P. Golle, editors). LNCS 5628:148–166, 2009.
12. The julia language. <http://julialang.org/> (visited 30/09/2015).
13. R.L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. Wiley, 1976.
14. P. Kumaraguru and L. F. Cranor. Privacy Indexes: A Survey of Westins Studies Technical Report, Institute for Software Research, Carnegie Mellon University. <http://www.cs.cmu.edu/~ponguru/CMU-ISRI-05-138.pdf>, 2005.
15. R. A. Posner. The right of privacy. *Georgia Law Review* 2 (3), 393–422, 1978.
16. R. A. Posner. The economics of privacy. *Amer. Econ. Rev.* 71 (2), 405–409, 1981.
17. Y.Pu and J.Grossklags. An Economic Model and Simulation Results of App Adoption Decisions on Networks with Interdependent Privacy Consequences. LNCS 8840, 246-265, 2014.
18. F. Roesner, B. T. Gill, and Tadayoshi Kohno Sex, Lies, or Kittens? Investigating the Use of Snapchat’s Self-Destructing Messages. In *Proc. Financial Crypto. and Data Security '14* N. Christin and R. Safavi-Naini, editors. LNCS 64–76, 2014.
19. B. Rossler. *The value of privacy*. Wiley, 2005.
20. G. J. Stigler. An introduction to privacy in economics and politics. (*The Journal of Legal Studies*) 9 (4), 623–644, 1980.
21. H. R. Varian. Economic aspects of personal privacy. In *Privacy and Self-regulation in the Information Age*. US Department of Commerce, 1997.
22. A. Westin. *Privacy and Freedom*. Atheneum Publishers, 1967.