COMPUTING SCIENCE

Formalization of Influencing in Information Security

Charles Morisset, Iryna Yevseyeva, Thomas Groß, and Aad van Moorsel

TECHNICAL REPORT SERIES

No. CS-TR-1423 May 2014

TECHNICAL REPORT SERIES

No. CS-TR-1423 May, 2014

Formalization of Influencing in Information Security

C. Morisset, I. Yevseyeva, T. Groß, and A. van Moorsel

Abstract

Information security decisions typically involve a trade-off between security and productivity. In practical settings it is often the human/user who is best positioned to make this trade-off decision, or in fact has a right to make its own decision (such as in the case of 'bring your own device'). It then may be useful to discuss approaches which aim to influence the user decision, while leaving end responsibility with the user. This is often referred to as nudging the user, or, more generally, as influencing human behavior. The main aim of this paper is to provide a generic formalization to facilitate rigorous quantitative analysis of influencing information security behavior, providing a theoretical basis for studying, optimizing, comparing and evaluating approaches. In particular, we propose an agent-based formalization that captures the human decision maker as well as the influencer and the relationship between them. Within this formalization we will characterize an optimal policy for influencing and formally prove that such policies are optimal. We then embed multi-criteria decision making into our formalism as an approach to model human behavior and to choose between alternatives. We apply our formalization by deriving optimal policies for the selection of WiFi networks, in which the graphical user interface aims to nudge the user to particular security behavior.

Printed and published by Newcastle University,
Computing Science, Claremont Tower, Claremont Road,
Newcastle upon Tyne, NE1 7RU, England.

Bibliographical details

MORISSET, C., YEVSEYEVA, I., GROß, T., VAN MOORSEL, A.

Formalization of Influencing in Information Security
[By] C. Morisset, I. Yevseyeva, T. Groß, and A. van Moorsel

Newcastle upon Tyne: Newcastle University: Computing Science, 2014.

(Newcastle University, Computing Science, Technical Report Series, No. CS-TR-1423)

Added entries

NEWCASTLE UNIVERSITY
Computing Science. Technical Report Series. CS-TR-1423

Abstract

Information security decisions typically involve a trade-off between security and productivity. In practical settings it is often the human/user who is best positioned to make this trade-off decision, or in fact has a right to make its own decision (such as in the case of 'bring your own device'). It then may be useful to discuss approaches which aim to influence the user decision, while leaving end responsibility with the user. This is often referred to as nudging the user, or, more generally, as influencing human behavior. The main aim of this paper is to provide a generic formalization to facilitate rigorous quantitative analysis of influencing information security behavior, providing a theoretical basis for studying, optimizing, comparing and evaluating approaches. In particular, we propose an agent-based formalization that captures the human decision maker as well as the influencer and the relationship between them. Within this formalization we will characterize an optimal policy for influencing and formally prove that such policies are optimal. We then embed multi-criteria decision making into our formalism as an approach to model human behavior and to choose between alternatives. We apply our formalization by deriving optimal policies for the selection of WiFi networks, in which the graphical user interface aims to nudge the user to particular security behavior.

About the authors

Charles Morisset is a Senior Research Associate at Newcastle University, working with Aad van Moorsel on quantitative aspects of security, in particular in the decision making process and in access control mechanisms. Charles received is PhD from Université Pierre et Marie Curie - Paris VI in France in 2007, on the topic of formalisation of access control systems. He then worked from 2007 to 2009 at the United Nations University, in Macau SAR, China, on formal methods for software engineering, after which he joined the Information Security Group at Royal Holloway, University of London, to work on risk-based access control until 2011. From 2011 to 2013, he worked at the Istituto di Informatica e Telematica in Pisa, Italy, on formal methods and access control, and he joined the Centre for Cybercrime and Computer Security at Newcastle University in 2013.

Iryna Yevseyeva is a research associate at the Choice Architecture for Information Security (ChAISe) project at Newcastle University. She contributes to the project with her expertise in optimisation and decision making, in particular, in multi-criteria optimisation and decision analysis. Before joining Newcastle University in 2013, she was a post-doctoral researcher on multi-objective optimisation: in the Netherlands the Leiden Institute of Advanced Computer Science at the Leiden University in 2012-2013 on drug discovery; and in Portugal at the Polytechnic Institute of Leiria in 2011-2012 on spam filtering; at INESC Porto with Ciencia 2008 grant in 2009-2011 on scheduling; at the University of Algarve with grants from Academy of Finland and European Commission (Erasmus Mundus) in 2008-2009 on algorithms development. Iryna received PhD degree in computer science and optimisation from the University of Jyvaskyla, Finland, in 2007, for the research on multi-criteria classification with applications in healthcare. Before joining PhD program in 2004, she worked as a software developer at the Niilo Maki Institute of Neuropsychology, Finland, in 2001-2003. She has Master of Science degree in mobile computing from the University of Jyvaskyla, Finland (2001) and Master degree with honours in information technology from the Kharkov National University of Radio-electronics, Ukraine (2000).

Thomas Groß is a tenured lecturer (assistant professor) in security, privacy and trust at the School of Computing Science at the University of Newcastle upon Tyne (since 2011). He is the director of the Centre for Cybercrime and Computer Security (CCCS), a UK Academic Centre of Excellence in Cyber Security Research (ACE-CSR). His research interests are in security and privacy as well as applied cryptography and formal methods. He was a tenured research scientist in the Security and Cryptography group of IBM Research - Zurich before that and IBM's Research Relationship Manager for privacy research. Thomas received his M.Sc. (Dipl. Inf.) in Computer Science at the Saarland University, Germany, in 2004. He received his Ph.D. (Dr.-Ing.) from the Ruhr-University Bochum, Germany, in 2009. His thesis was on the security analysis of standardized identity federation. Thomas is a member of the GI, ACM, IEEE, IACR and EATA, as well as Alumnus of the German National Academic Foundation.

Aad van Moorsel is a Professor in Distributed Systems and Head of School at the School of Computing Science in Newcastle University. His group conducts research in security, privacy and trust. Almost all of the group's research contains elements of quantification, be it through system measurement, predictive modelling or on-line adaptation.
Aad worked in industry from 1996 until 2003, first as a researcher at Bell Labs/Lucent Technologies in Murray Hill and then as a research manager at Hewlett-Packard Labs in Palo Alto, both in the United States. He got his PhD in computer science from Universiteit Twente in The Netherlands (1993) and has a Masters in mathematics from Universiteit Leiden, also in The Netherlands. After finishing his PhD he was a postdoc at the University of Illinois at Urbana-Champaign, Illinois, USA, for two years. Aad became the Head of the School of Computing Science in 2012.

Suggested keywords

DECISION-MAKING
NUDGING
UNCERTAINTY

Formalization of Influencing in Information Security

Charles Morisset, Iryna Yevseyeva, Thomas Groß, and Aad van Moorsel

Centre for Cybercrime and Computer Security
School of Computing Science, Newcastle University
Newcastle upon Tyne NE1 7RU, UK
firstname.lastname@newcastle.ac.uk

Abstract. Information security decisions typically involve a trade-off between security and productivity. In practical settings it is often the human/user who is best positioned to make this trade-off decision, or in fact has a right to make its own decision (such as in the case of 'bring your own device'). It then may be useful to discuss approaches which aim to influence the user decision, while leaving end responsibility with the user. This is often referred to as nudging the user, or, more generally, as influencing human behavior. The main aim of this paper is to provide a generic formalization to facilitate rigorous quantitative analysis of influencing information security behavior, providing a theoretical basis for studying, optimizing, comparing and evaluating approaches. In particular, we propose an agent-based formalization that captures the human decision maker as well as the influencer and the relationship between them. Within this formalization we will characterize an optimal policy for influencing and formally prove that such policies are optimal. We then embed multi-criteria decision making into our formalism as an approach to model human behavior and to choose between alternatives. We apply our formalization by deriving optimal policies for the selection of WiFi networks, in which the graphical user interface aims to nudge the user to particular security behavior.

1 Introduction

People continuously make information security decisions: should I use this wireless, should I put this person's USB stick in my laptop, how do I choose and memorize passwords? Almost always, the decision involves a trade-off between security and other concerns, such as being able to complete an important task or being able to easily do something that otherwise could be cumbersome. The decisions are often complex, with several objectives to be considered simultaneously, and the optimal decision may very much depend on the specific situation: while using a stranger's USB stick is not advisable, the importance of the job to be completed and/or knowledge about the owner of the USB stick may make it advisable to put the USB stick in one's laptop, despite the associated information security risks.

In situations such as above, a simple compliance policy (such as, not to allow USB sticks) would be suboptimal. Instead, one would want to allow some freedom for the owner of the laptop to decide the best course of action. In general terms, unless one can specify a compliance policy that is optimal under all possible circumstances, there is room for improvement by allowing the user to make the final decision. There exist other situations in which the user should play a role in the decision making. For instance, in case of BYOD (bring your own device) [5], where the device owner uses their own device for work-related activities, the fact that the user owns the device puts certain restrictions on what the employer can decide without the owner's input.

In all these situations the end user is involved in the information security decision making, and is in fact end responsible. Then, and this is key for this paper, it may be advisable that service providers (telecoms, online banks), device vendors, employers, or other parties are able to *influence* the decision making, without restricting the end user. In the literature this is often referred to as nudging [19]. Nudging leaves the choice with the user, but aims to influence the decision so that the user is more likely to make a beneficial decision, e.g., by presenting choices in a particular manner that aims to impact the choice a person ends up making. There are many aspects to nudging that deserve discussion, but in this paper we do not debate the specific approach, but aim to derive results for influencing in general.

To our knowledge there is no formal definition of nudge or similar concept of influencing exists, even though recently in [9], Heilmann presented schematically the nudge success conditions from perspective of influencing autonomous system, also called System 1 (and not reflective, also called System 2) [11], and showed the difference of these conditions for different types of nudges with respect to taxonomy of Bovens [3].

This paper provides a formalization of the concept of influencing that is as general as possible while assuring it is intuitive and useful. We believe such formalization is necessary to enable a solid quantitative treatment of optimization, comparison and evaluation of influencing approaches. In particular, we want to be able to apply mathematical optimization to decision making as well as to the decision on how to influence, and for that we need a rigorous underpinning and understanding of the problem at hand. We also want to understand how different variants of influencing behavior relate to each other, and the current literature [19] that introduces approaches in informal (albeit reasonably precise) language will always leave room for interpretation and misunderstanding. Finally, we want to be able to evaluate the level of success of influencing behaviors, be it experimentally or theoretically–again, a formal framework allows us to define the experimental or theoretical setting under which we carry out the evaluation. This paper will not reach all these goals, but provide the underlying quantitative framework, with an emphasis on the application of optimization techniques.

Our formalization follows naturally from the above scenarios: we present (in Section 2) an agent-based formalization, in which the decision maker and

influencer are agents that influence each other. We associate with agents (i) an observation by the agent of the environment that cannot be altered, although uncertainty about the environment can exist within the agent, and (ii) a context that is subject to attempts to be influenced. Given environment and context, a probabilistic policy is formulated that represents decisions (by the decision maker agent) as well as the approach to influencing (by the influencer agent); we introduce a notion of effect, which, up to a point, can modify decisions made by the decision maker. We prove that we can formally characterize the optimal policy using an impact function.

In Section 3 we embed Multi-Criteria Decision Making into our formalism, since it provides a natural way to discuss optimization of decisions, both for the decision maker and the influencer. We illustrate our formalism and its merits using a WiFi scenario taken from [20] in Section 4. In the WiFi example, a device user decides between networks and the device presents choices so as to influence the decision of the device user. In this case, the decision maker agent represents the device user and the influencer agent represents the way the device presents the choices. We will show optimal policies for the WiFi example and show that based on the uncertainty of the decision-maker as well as influencer, different effects can be optimal. In particular, we show that it is possible that it is better for the influencer *not* to attempt to influence the decision-maker, thus illustrating that there can be such a thing as 'too much security'. It also shows that our framework facilitates thinking about more flexible enforcement of security policies. Finally, Section 6 discusses possible extensions of the framework, in particular considering explicit time and multiple influencers.

2 Formalization of Influence

We consider here a multi-agent system, where each agent can make a decision at each step. An important characteristic of an agent is to have a partial view of the environment through its sensors, and on which it can have an impact through its actions. We first present a model considering a single decision-maker and a single influencer, in order to focus on the relationship between these two agents, and we propose in Section 5 different extensions of this basic model, including explicit times and multiple influencers.

2.1 Environment, Observation and Context

We write $\mathcal{E}$ for the set of possible environments, and given an agent a_i, we write Θ_i for the set of observations available to a_i, and $\Theta = \cup_i \Theta_i$ for the set of all possible observations. Intuitively, an observation represents a partial view of the environment, different for different agents. We encode the relationship between environments and observations through the probabilities $p(e \mid \theta)$ and $p(\theta \mid e)$, representing the conditional probabilities of being in the environment e when observing θ and of observing θ when being in the environment e.

We also write $\mathcal{C}_i$ for the set of contexts available to a_i and $\mathcal{C} = \cup_i \mathcal{C}_i$ for the set of all contexts. Intuitively, an observation corresponds to a probe from real-world information, and as such, cannot be changed, but can be uncertain, while a context is defined within the agent, and as such can be modified, as described later in this section. The separation between observation and context is part of the modelling process, and can change from one system to another, depending on the assumptions considered. In the WiFi scenario, an observation can for instance include the strength of the signal or the level of trust of a network (which can be uncertain), while a context might include the color in which the networks are displayed.

2.2 Decision-Maker

A decision-maker is an agent responsible for selecting an alternative, given an observation and a context. We write $\mathcal{A}_i$ for the set of alternatives available to a_i, such that $\mathcal{A}_i \neq \emptyset$, meaning that an agent has always at least one alternative possible (such an alternative could be reduced to a skip action, i.e., an action with no impact on the environment).

A decision-maker is a *probabilistic* agent, thus allowing for the modelling of groups of users rather than single users. More formally, a decision-maker is an agent a_i associated with a policy[1] $\pi_i : \Theta_i \times \mathcal{C}_i \to \mathcal{P}(\mathcal{A}_i)$.

We say that a decision-maker is deterministic given an observation θ_i and a context c_i whenever there is an alternative b such that[2] $\pi_i(\theta_i, c_i, b) = 1$. In this case, we abuse the notation and write $\pi_i(\theta_i, c_i) = b$.

2.3 Influencer

An *effect* is a unary function that can modify the context of an agent. Given an agent a_i, we say that an effect applicable to a_i is a function $\eta : \mathcal{C}_i \to \mathcal{C}_i$, and we write $\mathcal{N}_i$ for the set of effects applicable to a_i. For the sake of generality, we assume that $\mathcal{N}_i$ is closed under composition, i.e., if $\eta \in \mathcal{N}_i$ and $\eta' \in \mathcal{N}_i$, then $\eta \circ \eta' \in \mathcal{N}_i$.

An influencer is a deterministic agent (as we will explain) responsible for selecting an effect over a_i at each time. Such an agent does not use any context to choose an effect, preventing it by construction to be influenced. However, as we discuss in Section 5, we do not see this as a limitation, but rather as a design choice, in order to avoid influencing loops. In addition, without loss of generality, we assume that an influencer can only influence a single agent, since influencing multiple agents can be simulated by several influencers, one for each influenced agent.

---

[1] Given a set X, we write $\mathcal{P}(X)$ for the probability space associated with X, i.e., for the set of functions $f : X \to [0, 1]$ such that $\sum_{x \in X} f(x) = 1$

[2] Strictly speaking, we should write $\pi_i(\theta_i, c_i)(b)$, however, for the sake of clarity, we use here the curried notation for probability functions, and we write $\pi_i(\theta_i, c_i, b)$ instead.

The motivation of this approach is to analyse how an agent behaves when influenced, and to evaluate whether this influence is beneficial or not. Hence, although a decision-maker can be probabilistic, denoting a possible behaviour of the agent, the application of effects should somehow be deterministic, so that given one decision-maker, there is a single agent corresponding to the influenced decision-maker. Hence, an influencer is not a probabilistic agent, and deterministically selects an effect (that can be composed of several effects, or reduced to the identity function) to apply at each time.

More formally, an influencer over a decision-maker a_i is an agent a_j associated with a policy $\pi_j : \Theta_j \to \mathcal{A}_j$, such that $\mathcal{A}_j \subseteq \mathcal{N}_i$. In this case, we write $a_j \rhd a_i$. We also write $\pi_j(\theta_j)$ to denote the specific effect η_j when agent a_j observes θ_j.

Note that $\mathcal{A}_i$ does not change with the application of effects. In other words, influencing an agent does not change the set of alternatives possible for that agent. However, an effect can change the probability of an alternative to change from 0 to a value strictly positive, which in practice means that this alternative had no chances to be selected before the effect, and is possible after. Hence, although we do not consider here the dynamic creation/suppression of alternatives, our approach is flexible approach to consider the dynamic activation/deactivation of alternatives.

2.4 Influence

The raw impact of an influencer can be measured in a differential way: given an observation θ_j for a_j, we say that a_j has an impact on a_i whenever $\pi_i(\theta_i, \eta_j(c_i)) \neq \pi_i(\theta_i, c_i)$. In other words, a_j has an impact on a_i if and only if a_i would behave differently without a_j.

The above only measures a difference in behaviour, but does not indicate whether the influence is in the direction desired by a_j. Hence, we introduce a *impact function* $\rho_i : \mathcal{E} \times \mathcal{A}_i \to \mathbb{R}$ such that, given an environment e and an alternative b, $\rho_i(e, b)$ represents the impact of a_i selecting b. In the rest of the paper, we consider that the impact function intuitively represents a cost on the system, and as such, the aim of an influencer is to minimize the impact, i.e., a lower impact is 'better'. Note, the impact function should be seen as an ideal valuation of the possible alternatives, and as a way to evaluate the behaviour of the agents, rather than as a way to define the behaviour of the agents.

Definition 1. *Given a decision-maker a_i with an observation θ_i and a context c_i, the impact on an environment e of an effect η is given by $\delta_i(\theta_i, e, \eta)$, defined as follows:*

$$\delta_i(\theta_i, e, \eta) = \sum_{b \in \mathcal{A}_i} \pi_i(\theta_i, \eta(c_i), b)\, \rho_i(e, b)$$

We can now define the *global impact* of an influencer in a given environment, which depends on the probabilities of the observations made by the influencer and the decision maker.

Definition 2. *Given an environment e, the global impact of an influencing policy π_j from the agent a_j over a_i is defined as:*

$$\Delta_i^{\pi_j}(e) = \sum_{\theta_j \in \Theta_j} p(\theta_j \mid e) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j(\theta_j))$$

Note that Definition 2 assumes that for any given environment, the agents make statistically independent observations (of course, this does not imply different or arbitrary observations are made, but only implies that $p(\theta_i, \theta_j \mid e) = p(\theta_i \mid e)p(\theta_j \mid e)$). This is justified by the design of the formalism, in which dependence between agents is through context, not through observations.

2.5 Optimal Influencer

We say that a policy for an influencer is *optimal* whenever it minimizes the expected impact of the influencer.

Definition 3. *A policy π_j for an influencer a_j is optimal if, and only if, for any other policy π'_j:*

$$\sum_{e \in \mathcal{E}} p(e) \Delta_i^{\pi_j}(e) \leq \sum_{e \in \mathcal{E}} p(e) \Delta_i^{\pi'_j}(e)$$

In an ideal setting, the influencer knows both the environment e and the observation made by the decision-maker θ_i, and so can simply select the effect η that minimizes $\delta_i(\theta_i, e, \eta)$.

Proposition 1 *The optimal policy for a_j such that $a_j \rhd a_i$ is defined by, for any θ_j:*

$$\pi_j^*(\theta_j) = \arg\min_{\eta \in \mathcal{A}_j} \left[\sum_{e \in \mathcal{E}} p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \eta) \right]$$

The proof can be found in Appendix A, and roughly speaking, relies on the application of Bayes' Theorem over $p(e) \cdot p(\theta_j \mid e)$.

2.6 Extreme Influencers

In practice, the probabilities of $p(e \mid \theta_j)$ and $p(\theta_i \mid e)$ might not be known by a_j, in which case the policy π_j^* cannot be computed. In this case, the influencer can only make a decision based on its own observation θ_j of the environment. Let us consider two extreme strategies: non-influence and forced influence.

We write id for the neutral effect (i.e., $id(c) = c$, for any context c) and π_{id} for the policy such that $\pi_{id}(\theta_j) = id$. The impact of this policy is given by:

$$\Delta_i^{\pi_{id}}(e) = \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e) \sum_{b \in \mathcal{A}_i} \pi_i(\theta_i, c, b)\, \rho_i(e, b)$$

We say that a policy π_j is *beneficial* for a_i if, and only if, $\sum_{e\in\mathcal{E}} p(e)\Delta_i^{\pi_j}(e)$ is lower than $\sum_{e\in\mathcal{E}} p(e)\Delta_i^{\pi_{id}}(e)$. In other words, an influencer is beneficial if, and only if, its impact is better than applying no effect. It is worth observing that as long as the effect *id* is available to an influencer for all observations, then the optimal strategy is always beneficial.

Another extreme strategy is for the influencer to *force* the decision-maker towards a specific alternative. Let us first assume that for each alternative $b \in \mathcal{A}_i$, there exists an effect η_b such that $\pi_i(\theta_i, \eta_b(c), b) = 1$, for any observation θ_i and any context c. In other words, η_b forces a_i to select b. A forcing influence for an influencer a_j could then be to select the effect forcing towards the alternative it believes to be the best.

Let us now assume that a_j knows the probability function $p(e \mid \theta_j)$, then we can define the forced influence policy:

$$\pi_F(\theta_j) = \varphi_{b_{\min}} \text{ where } b_{\min} = \arg\min_{b\in\mathcal{A}_i} \sum_{e\in\mathcal{E}} p(e \mid \theta_j) \cdot \rho_i(e, b)$$

and it follows that the impact of this policy is given by, for any environment e:

$$\Delta_i^{\pi_F}(e) = \sum_{\theta_j\in\Theta_j} p(\theta_j \mid e) \left[\min_{b\in\mathcal{A}_i} \sum_{e'\in\mathcal{E}} p(e' \mid \theta_j)\, \rho_i(e', b)\right]$$

It is worth observing that when a_j has a perfect knowledge of the environment, i.e., when given an observation θ_j, there is a single environment e such that such that $p(e \mid \theta_j) = 1$, then $\Delta_i^{\pi_F}(e) = \min_{b\in\mathcal{A}_i} \rho_i(e, b)$ is clearly minimal. In other words, when the influencer has perfect knowledge of the environment and has the possibility of forcing the decision-maker to select a given alternative, then doing so is optimal. Conversely, when either of the above conditions is not met, the traditional approach of forcing the decision maker to follow a particular alternative is not necessarily optimal.

3 Analysis of the Alternatives

In order to model human decision making and to evaluate the different alternatives for a decision-maker, we consider Multi Criteria Decision Analysis (MCDA). MCDA is particularly useful in situations where alternatives are evaluated on multiple, often conflicting, criteria, in search of solutions that represent the best trade-off(s) between these criteria. In information security, this trade-off is usually between security and productivity/usability, for instance, deciding between a more secure network and a faster one.

We first present the basics of MCDA and Multi-Attribute Utility Theory (MAUT). We then detail how to define the impact function and the policy of a decision maker, an d finally, we illustrate this approach for the WiFi example.

3.1 Multi-Attribute Utility Theory

Multi-attribute utility theory [13] is an MCDA approach that assumes that decision makers maximize their implicit utility function. MAUT is a compensatory technique, since it allows smaller values on a subset of criteria to be compensated by a large value on a single criterion, and is based on expected utility theory with some strong technical assumptions related to comparability, transitivity, continuity, and independence of outcomes (that assumes independence of criteria). MAUT is attractive because of its sound theoretical foundations (based on expected utility theory), its non-monetary nature and as a basis for comparison of new, not yet considered alternatives with the same utility function constructed for the same decision maker. In addition, its natural approach to modeling risk behavior is particularly attractive for designing security decisions, where risk attitude of decision makers plays crucial role in their decision patterns.

In MCDA alternatives are evaluated and compared using a set of criteria $\mathcal{G}$, such that each criterion should be either minimized or maximized (the direction of optimization). Each criterion comes with a scale, in which alternatives can be compared. Typical scales include real numbers, intervals, ratios, binary or verbal values (qualitative descriptions), which are ordered with respect to the optimization direction. Each criterion $g \in \mathcal{G}$ is therefore associated with a scale $\mathcal{K}_g$, and we write $g_{\min} \in \mathcal{K}_g$ and $g_{\max} \in \mathcal{K}_g$ for the minimal and maximal values of g. We write $\mathcal{K} = \bigcup_{g \in \mathcal{G}} \mathcal{K}_g$ for the set of all possible scales, and without any loss of generality, we assume that all criteria are maximized (minimized criteria can be simply multiplied by -1).

The values of alternatives for different criteria are provided through observations, environments and contexts. Since decision-makers and influencers use different sources for these, for the sake of generality, we introduce an abstract set of states $\mathcal{S}$, such that, for any $\sigma \in \mathcal{S}$, $\sigma(g, b)$ denotes the value of the alternative b for the criterion b.

The global value of an alternative is obtained by aggregating the value for all criteria. However, before aggregation, these values must be *normalized*, in order to compare comparable values. A normalization function, which in MAUT corresponds to *marginal utility function*, is a function $n : \mathcal{G} \times \mathcal{K} \to [0, 1]$. This function can change from one decision-maker to another, thus encoding some notion of preference.

In addition, preferences can be encoded using criteria weights, which in MAUT represent trade-offs between criteria. Here, a weight shows the relative importance of the criterion, when compared to other criteria. In particular, it defines how many units of one criterion can be traded-off for a unit of another criterion. The criteria weights can change depending on the state, and therefore we define the function $w : \mathcal{S} \to \mathcal{P}(\mathcal{G})$. We can now define the notion of MAUT model.

Definition 4. *A MAUT model is a tuple $M = (\mathcal{A}, \mathcal{G}, \mathcal{S}, n, w)$, where $\mathcal{A}$ is a set of alternatives, $\mathcal{G}$ a set of criteria, $\mathcal{S}$ a set of states, $n : \mathcal{G} \times \mathcal{K} \to [0, 1]$ is a normalisation function and $w : \mathcal{S} \to \mathcal{P}(\mathcal{G})$ is a weight function.*

After mapping all criteria utilities to their scales, normalizing them and defining weights, the alternatives can be evaluated. For aggregating marginal criteria utilities for each alternative some form of aggregation function should be used, e.g. multiplicative, additive or some combination of both is usually applied. We will explain in Section 4.2 the relation between the criteria value and the impact function that we used in our framework. For now, we introduce one of simplest forms of aggregating evaluations on all criteria scores for each alternative is weighted sum, which we will use in the WiFi example:

Definition 5 (Utility function). *Given a model $M = (\mathcal{A}, \mathcal{G}, \mathcal{S}, n, w)$, the utility of an alternative $b \in \mathcal{A}$ given a state $\sigma \in \mathcal{S}$ is defined as:*

$$v(b, \sigma) = \sum_{g \in \mathcal{G}} w(\sigma, g) \cdot n(g, \sigma(g, b)).$$

3.2 Decision Maker

We assume that decision makers are rational, and that they base their decision-making process using a MAUT model. Following the separation between observation and context, we consider two distinct sets of criteria, $\mathcal{G}_i^{\mathcal{O}}$ and $\mathcal{G}_i^{\mathcal{C}}$. Furthermore, a state consists of the union of the observation θ_i and of the context c_i, such that, intuitively, θ_i provides the value for the criteria in $\mathcal{G}_i^{\mathcal{O}}$ while c_i provides the value for the criteria in $\mathcal{G}_i^{\mathcal{C}}$. Hence, the set of states is defined as $\mathcal{S}_i = \Theta_i \cup \mathcal{C}_i$.

The model of a deterministic decision-maker a_i is therefore a tuple $M_i = (\mathcal{A}_i, \mathcal{G}_i^{\mathcal{O}} \cup \mathcal{G}_i^{\mathcal{C}}, \Theta_i \cup \mathcal{C}_i, n_i, w_i)$, and given an observation θ_i and a context c_i, the policy of a_i is defined as:

$$\pi_i(\theta_i, c_i) = \arg\max_{b \in \mathcal{A}_i} v(b, \theta_i \cup c_i).$$

Note that in order for a_i to be deterministic, we assume the existence of an arbitrary ordering over alternatives, so that if there are several alternatives maximizing the value function, a_i selects the highest one according to that ordering.

In Section 2, we defined decision-makers to be probabilistic in order to account for a group of users instead of a single one. We model this aspect in MAUT by considering a set $W \subseteq (\mathcal{S} \to \mathcal{P}(\mathcal{G}))$ equipped with a probability distribution $\psi : \mathcal{P}(W)$, such that, given a weight function $w \in W$, $\psi(w)$ represents the probability of w. We write π_w to represent the policy of the deterministic agent defined according to the weight function w, and the policy π_W of the population represented by (W, ψ) is defined as:

$$\pi_W(\theta_i, c_i, b) = \sum_{w \in W} \{\psi(w) \mid \pi_w(\theta_i, c_i) = b\} \tag{1}$$

3.3 Impact Function

The impact function ρ_i for a decision-maker a_i defines the impact of each alternative in an environment. In general, this function can be defined in many different ways (for instance, through an access control policy stating which alternatives are secure). We propose here to define it using a MAUT model, which is however slightly different from the one defined above. Indeed, an important aspect of the impact function is that it does not take the context into account.

Hence, given a set of alternatives $\mathcal{A}_i$, a set of criteria $\mathcal{G}_i^{\Theta}$, a set of environments $\mathcal{E}$ (providing the values for the criteria in $\mathcal{G}_i^{\Theta}$), a normalization function n and a weight function w, we can define the model $\overline{M_i} = (\mathcal{A}_i, \mathcal{G}_i^{\Theta}, \mathcal{E}, n, w)$. Note that the functions n and w need not be equal to those of the decision-maker a_i, and in practice, such a difference can reflect a misalignment between the preferences of the decision-maker and those of the influencer.

The impact function can then be defined directly as the utility function $\overline{v}$ of $\overline{M_i}$. However, in the context of information security, we want to clearly distinguish the alternatives, so that there "good" and "bad" alternatives. Hence, we define the impact function as, given an environment e and an alternative b:

$$\rho_i(e, b) = \begin{cases} 0 & \text{if } \bar{v}(b, e) \geq \bar{v}(b', e), \text{ for any } b', \\ 1 & \text{otherwise.} \end{cases}$$

In other words, an alternative has no impact if, and only if, it is maximal according to the utility function.

4 Case study: Selection of a WiFi Network

As a case study, let us consider an example of nudging in choosing a wireless network (WiFi). This is of particular interest in the context of BYOD, since employees work on their own devices and define security protection of their devices by themselves, thus potentially exposing sensitive information [18]. The dangers of choosing non-secure Wi-Fi are well documented: it exposes device and data transmitted to increased chances of spoofing and man-in-the-middle attacks [1], [16]. In general, over one billion workers will work remotely by 2015, over a third of the total worldwide workforce [10]. A company that allows BYOD may want to influence the employee so that the trade-off decision between security and productivity is done in the company's interest. Alternatively, users may want to have nudging software on their phone to assist in making the information security decisions for work as well as home use.

We introduced nudging in the security context of WiFi selection in [20]. There the focus was on introducing the user interface design nudges, and evaluate them with a user group. We consider here the traffic light effect [6], also used in framing choices [4], which assumes that users tend to select items identified with the green color, reject those identified with the red color, and do not differentiate when the color is neutral.

Table 1. Decision matrix for $\theta_1 = [s \mapsto (1,1), f \mapsto (0,2)], c_1 = [s \mapsto N, f \mapsto N]$

	criteria			
	trust	strength	color	
	$\{0,1,2\}$	$\{0,1,2\}$	$\{R,N,G\}$	(scale)
	$t \to \max$	$r \to \max$	$l \to \max$	(direction)
	0.5	0.3	0.2	(weights)
alternative				
s	1	1	N	
f	0	2	N	
$\pi_i(\theta_1, c_1)$		f		

4.1 Decision-maker

Let us consider a user a_i in a coffee-shop having to choose between two different networks $\mathcal{A} = \{s, f\}$: s is a secure network with weak signal; f is a public WiFi of the coffee shop, with strong signal, but not necessary safe. We want to illustrate with this example the trade-off between security and productivity/usability, and therefore consider the set of criteria $\mathcal{G}_i^{\mathcal{O}} = \{t, r\}$, indicating the trust and the strength of the network, respectively, and for the sake of simplicity, we assume that the scales are defined as $\mathcal{K}_t = \mathcal{K}_r = \{0,1,2\}$ (the higher the better).

Note that we consider here a simple and abstract notion of trust, and in practice, this notion can be defined using the presence in the white list of WiFi network providers predefined by security officer or system administrator of the company or by the employee itself. More sophisticated evaluation of 'trust' criterion may take into account current location of an employee [7].

We also consider a single context criterion $\mathcal{G}_i^{\mathcal{C}} = \{l\}$, which indicates the color in which the name of the network is displayed, with a scale $\mathcal{K}_l = \{R, N, G\}$, corresponding to red, neutral and green, respectively, and with the values 0, 0.5 and 1, respectively.

Finally, let us consider that the decision maker has the following weight function, for any observation θ and any context c: $w((\theta, c), t) = 0.5$, $w((\theta, c), r) = 0.3$ and $w((\theta, c), l) = 0.2$, meaning that connecting to a trusted WiFi is slightly more important for the decision maker than choosing WiFi with strong signal. The color of the presented name of a WiFi is less significant for the decision maker. We also consider that a_i uses linear normalization:

$$n_i(g, \theta, c) = \frac{g - g^{min}}{g^{max} - g^{min}}. \tag{2}$$

We use an extensional notation for observations, and we write $\theta = [s \mapsto (v_1, v_2), f \mapsto (v_3, v_4)]$ for the observation θ associating s with a trust of v_1 and strength of v_2, and f with a trust of v_3 and a strength of v_4. We use a similar notation for contexts. Table 1 represents the traditional decision matrix

[2] for a_i, for the observation $\theta_1 = [s \mapsto (1,1), f \mapsto (0,2)]$, and the context $c_1 = [s \mapsto N, f \mapsto N]$. We can therefore calculate the utility for each alternatives:

$$v(s, \theta_1, c_1) = 0.5 * 0.5 + 0.3 * 0.5 + 0.2 * 0.5 = 0.5$$
$$v(f, \theta_1, c_1) = 0.5 * 0 + 0.3 * 1 + 0.2 * 0.5 = 0.4$$

and it follows that a_i selects $s = \pi_i(\theta_1, c_1)$.

It is worth observing that the slight difference in the observation can lead to a different final choice. For instance, consider the previous observation where the trust of f is now equal to 1 (instead of 0), encoded by the observation $\theta_2 = [s \mapsto (1,1), f \mapsto (1,2)]$. The utilities of the alternatives become $v(s, \theta_2, c_1) = 0.5$ and $v(f, \theta_2, c_1) = 0.65$, and therefore a_i now selects $f = \pi_i(\theta_2, c_1)$.

Finally, we can model a population of users using a set of weight functions, instead of a single one. For instance, consider the set $W = \{w_1, w_2, w_3\}$ where, for any observation θ and any context c: $w_1((\theta, c), t) = 0.3$, $w_1((\theta, c), r) = 0.6$ and $w_1((\theta, c), l) = 0.1$; $w_2((\theta, c), t) = 0.5$, $w_2((\theta, c), r) = 0.4$, and $w_2((\theta, c), l) = 0.1$; and $w_3((\theta, c), t) = 0.3$, $w_3((\theta, c), r) = 0.3$ and $w_3((\theta, c), l) = 0.4$. Let us consider a probability function ψ such that $\psi(w_1) = \psi(w_2) = \psi(w_3) = 1/3$. Then, we can observe that $\pi_{w_1}(\theta_1, c_1) = f$, $\pi_{w_2}(\theta_1, c_1) = s$ and $\pi_{w_3}(\theta_1, c_1) = s$ (the utility of s and f are equal with w_3, and we arbitrarily choose s), and therefore, following Equation 1, we have $\pi_W(\theta_1, c_1, f) = 1/3$ and $\pi_W(\theta_1, c_1, s) = 2/3$.

4.2 Influencer

As described above, we consider a single context criterion l, corresponding to the color in which a network is displayed. Hence, by construction, the set of effects is limited to effects of the form η_{xy}, such that, given any context c, $\eta_{xy}(c) = [s \mapsto x, f \mapsto y]$.

From the example in the previous section, let us consider again the first observation θ_1, where a_i selects s with the neutral context c_1, and let us now consider the context $c_2 = \eta_{RG}(c_1)$. The utilities of the alternatives become $v(s, \theta_1, c_2) = 0.4$ and $v(f, \theta_1, c_2) = 0.5$, leading a_i to select f: $\pi_i(\theta_1, c_2) = f$. In other words, by changing the color of f to green and that of s to red, the influencer managed to change the decision of the decision-maker from s to f. However, note the impact of this effect depends on the environment. For instance, if the utility of f are null and that of s are maximal, there is no effect that will make a decision-maker with a weight on $l \neq 1$ change its decision from s to f. Similarly, if the decision-maker has a weight equal to 0 on the criterion l, then all effects have no impact.

4.3 Optimal Effect

We now illustrate the definition of the optimal policy π^*_j, and of the corresponding optimal effect η^*_{xy}. Let us define the environments $e_1 = [s \mapsto (1,1); f \mapsto (0,2)]$ or $e_2 = [s \mapsto (1,1); f \mapsto (1,2)]$ (i.e., respectively corresponding to the observations θ_1 and θ_2 defined above). We assume that the influencer a_j is

able to measure a probability of each environment from the set of all possible ones to take place $e_1, e_2 \in \mathcal{E}$, when observing θ_j: $p(e_1 \mid \theta_j = \theta_1) = 0.7$ and $p(e_2 \mid \theta_j = \theta_2) = 0.3$, respectively.

We consider here the case where the influencer worries more about the trust of connection rather than the strength of signal, and so the weight function for the utility function of the impact function is defined as $w(e,t) = 0.6$ and $w(e,r) = 0.4$, for any environment e. It follows that $\rho_i(e_1, s) = 0$ and $\rho_i(e_1, f) = 1$, while $\rho_i(e_2, s) = 1$ $\rho_i(e_2, f) = 0$. In other words s and f are the alternatives preferred by the influencer in e_1 and e_2, respectively. Let us finally consider the probabilities associated with the observations of the decision-maker. We assume that in e_1, a_i is certain to observe θ_1, while in e_2, a_i can observe either θ_1 or θ_2. More precisely, we define $p(\theta_i = \theta_1 \mid e_1) = 1$, $p(\theta_i = \theta_1 \mid e_2) = 0.2$ and $p(\theta_i = \theta_2 \mid e_2) = 0.8$.

To select an optimal effect η^*_{xy}, we have to compute the impact of each effect, and select the one with the minimal impact. For instance, given the decision maker described above, i.e., associating a weight of 0.3, 0.5 and 0.2 to the trust, the strength and the color of a network, respectively, we can calculate that for any context c, $\pi_j(\theta_1, \eta_{RG}(c)) = f$, while $\pi_j(\theta_1, \eta_{xy}(c)) = s$ for any $\eta_{xy} \neq \eta_{RG}$ (details can be found in Appendix B, Table 2). Hence, we have that $\delta_i(\theta_1, e_1, \eta_{RG}) = 1$, and $\delta_i(\theta_1, e_1, \eta_{xy}) = 0$ for any $\eta_{xy} \neq \eta_{RG}$. In other words, if the influencer knows that the environment e_1, then any effect but η_{RG} is optimal.

We can similarly compute $\delta_i(\theta_1, e_2, \eta)$ equals 0 if $\eta = \eta_{RG}$ and equals 1 otherwise, (i.e, the case dual to θ_1 and e_1) and $\delta_i(\theta_2, e_2, \eta)$ equals 1 if $\eta = \eta_{GR}$ and equals 0 otherwise (details can be found in Appendix B, Table 2 and 3). We can then compute the global impact of each effect given the initial observation of the influencer. Writing $\gamma_j(\theta_j, \eta) = \sum_{e \in \mathcal{E}} p(e|\theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i|e)\delta_i(\theta_i, e, \eta)$, we have:

$$\gamma_j(\theta_1, \eta) = \begin{cases} 0.7 & \text{if } \eta = \eta_{RG}, \\ 0.3 & \text{if } \eta = \eta_{GR}, \\ 0.06 & \text{otherwise.} \end{cases}$$

In other words, there is evidence against applying strong effects η_{RG} and η_{GR} that 'swap' the choice of the decision maker, since they can have a bad impact when applied wrongly, while the other effects are all equally optimal.

This example shows that small differences applied to changing presentation of each alternative does not necessarily change the decision-maker selection strategy, i.e., changing the color of only one network might not have any impact. Furthermore, and perhaps one illustration of the interest of influence with respect to traditional enforcement mechanisms, "doing nothing" (or presenting both alternatives in the same or neutral color) can be one of the best options in these cases. Clearly, this situation is due to the uncertainty relative to the influencer, which is not able to capture precisely the impact of its influence.

5 Extensions

The model we have presented so far does not depend on explicit time, and is limited to a single influencer per decision-maker. Although this model was enough to express the core contribution of our approach, which is the quantification of the flexible enforcement of security policies, following nudging techniques, more complex models might be required. We now present two extensions, the first introducing time and the second allowing multiple influencers.

System Evolution In several cases, the observation made by a decision-maker or an influencer at time t depends on the decisions made at time $t' < t$. For instance, in the WiFi example, if one network is not encrypted, selecting it might lead to the decision-maker to decide to set up its own encryption system.

In addition, the effects applied on the context at time t might persist at time $t' > t$. In this case, we could define $c_i^{t+1} = \eta_i(c_i^t)$, where c_i^t denotes the context of the agent a_i at time t. In general, given an agent a_i influenced by a_j we have the following dependencies at time t: (i) The observations θ_i^t and θ_j^t depend on the environment e^t and can depend on decisions made by all agents (not restricted to a_i and a_j) in the system at time $t' < t$; (ii) The decision made by a_j depends on the observation θ_j^t; (iii) The context c_i^t depends on the context c_i^{t-1} and the decision $\pi_j^t(\theta_j^t)$ made by a_j at time t (and on the decision made by other influencers if any); (iv) The decision made by a_i depends on θ_i^t and c_i^t.

In order to avoid loops in the decision process, the environment e^t should not depend on decision made at time t, but can however depend on decisions made at time $t' < t$.

Multiple Influencers In practice, a single decision-maker is influenced by multiple influencers. For instance, the employee of a company can be influenced both by the business department, in order to increase productivity, and by the technical department, in order to increase security.

However, several difficulties arise from the consideration of multiple influencers. The first one is the ordering of the effect application. Indeed, the impact of applying different effects coming from different influencers might be different according to the order in which they are applied. For instance, consider the WiFi effects $\eta_{ll'}$ as described in Section 4: clearly, the last effect applied takes over all effects previously applied. In order to address this problem, we could only consider commutative effects: given $\eta, \eta' \in \mathcal{N}_i$, $\eta \circ \eta' = \eta' \circ \eta$. A more general approach is to define an *effect scheduler*, which collects the effects from the different influencers, and applies them on the context.

The second difficulty is that of influencing loops. Indeed, due to the synchronicity of the application of effects, a_j and a_k could influence each other mutually, potentially indirectly, which means that the alternative chosen by a_j depends on that chosen by a_k, which itself depends on that chosen by a_j, thus creating a loop. Note that strictly speaking, since an influencer does not use a context, it cannot be influenced, so as long as the observation at a given time

do not depend on the decisions made at the same time, there cannot be such loops. However, being able to characterise a chain of influence seems appealing to model some real-world scenarios (for instance, the influences that can come through the different levels of management in a company), in which case a possible solution is to delay the application of an effect, such that the effect chosen by a_j at time t is only applied to the context of a_i at $t+1$.

Note that a particularly interesting case for considering multiple influencers is when one of the influencers is a malicious agent. In this case, the motivation of an influencer a_j can be to counter the influence of another agent over a_i. In this case, a_j can aim at selecting the inverse effect to that chosen by the malicious agent.

6 Conclusion

In this work we have proposed a formalization of the concept of influence in information security systems based on a multi-agents approach. We have illustrated the approach with an example of selection a public WiFi, and have shown that due to uncertainty when making such security decisions, different effects may be optimal. In particular, in some cases not influencing will be the best strategy to follow, since influencing towards the wrong alternative can be counter-productive.

We believe we have made a first step towards adapting multi-attribute utility theory to deviations of human behavior from the rational one currently studied in behavioral economics and organizational psychology, see e.g. [12] and [11]. Further extensions can be considered, such as the integration of expected uncertain utility theory [8], which operates on interval utilities when modelling decision maker preferences. It also differentiates between attitude towards and perception of risk, which is particularly important for security decisions. Taking into account complexity of different criteria, such as 'trust' criterion, MAUT contribution may be further investigated by modeling more complex shapes of marginal utility functions, such as convex or concave utility functions corresponding to risk attitude (risk-prone or risk-averse) of decision maker, when compared to the linear marginal utility functions modelled here. Finally, the quantities obtained through MAUT can be used to characterize the strength of the effect applied, following for instance our recent approach in the context of quantitative access control policies [15].

In addition to the color effect, there is a large experience of behavioral sciences when developing and implementing security products revealed recently [14], [17], and other effects should be studied in the context of framework. For instance, in [20], in addition to the color, the ordering of the different networks presented to the decision maker by default is shown to have an impact.

Finally, the WiFi scenario provides an interesting basis for future work. For instance, in [7], the importance of name when choosing a WiFi was studied in the context of trust. Actually, the 'trust' criterion is interesting in itself, as it may

take into account various information, e.g. about the current location of the decision maker to avoid situations, where the most trusted network for a researcher located in a commercial center far away from university campuses appears to be the 'eduroam' WiFi, an internationally available network for all researchers and academic staff of universities provided within campuses of universities only.

References

1. M. Aime, G. Calandriello, and A. Lioy. Dependability in wireless networks: Can we rely on WiFi? *Security Privacy, IEEE*, 5(1):23–29, Jan 2007.
2. V. Belton and T. Stewart. *Multiple Criteria Decision Analysis: An Integrated Approach*. Kluwer Academic Publishers, Dordrecht, 2002.
3. L. Bovens. The ethics of nudge. In T. Grne-Yanoff and S. Hansson, editors, *Preference Change: Approaches from Philosophy, Economics and Psychology*, pages 207–219. Springer, Theory and Decision Library A, Berlin and New York, 2008.
4. E. K. Choe, J. Jung, B. Lee, and K. Fisher. Nudging people away from privacy-invasive mobile apps through visual framing. In *INTERACT (3)*, volume 8119 of *LNCS*, pages 74–91. Springer, 2013.
5. J. Clarke, M. G. Hidalgo, A. Lioy, M. Petkovic, C. Vishik, and J. Ward. Consumerization of it: Top risks and opportunities, 2012.
6. G. Farnham and K. Leune. Tools and standards for cyber threat intelligence projects, 2013.
7. A. Ferreira, J.-L. Huynen, V. Koenig, G. Lenzini, and S. Rivas. Socio-technical study on the effect of trust and context when choosing WiFi names. In *STM*, volume 8203 of *LNCS*, pages 131–143. Springer Berlin Heidelberg, 2013.
8. F. Gul and W. Pesendorfer. Expected uncertain utility theory. *Econometrica*, 82(1):1–39, 2014.
9. C. Heilmann. Success conditions for nudges: a methodological critique of libertarian paternalism. *European Journal for Philosophy of Science*, 4(1):75–94, 2014.
10. Aidc worldwide mobile worker population 2010-2015 forecast, 2012.
11. D. Kahneman. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York, 2011.
12. D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
13. R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. J. Wiley, New York, 1976.
14. I. J. Martinez-Moyano, S. H. Conrad, and D. F. Andersen. Modeling behavioral considerations related to information security. *Computers & Security*, 30(67):397 – 409, 2011.
15. C. Morisset, T. Gross, A. van Moorsel, and I. Yevseyeva. Nudging for quantitative access control systems. In *Human Aspects of Information Security, Privacy and Trust, HCII*. Springer, 2014. to appear.
16. MWRInfoSecurity. Mobile devices, 2013.
17. S. L. Pfleeger and D. D. Caputo. Leveraging behavioral science to mitigate cyber security risk. *Computers & Security*, 31(4):597–611, 2012.
18. J.-M. Seigneur, P. Kölndorfer, M. Busch, and C. Hochleitner. A survey of trust and risk metrics for a byod mobile worker world. In *Proceedings of SOTICS 2013*, pages 82 – 91. IARIA, 2013.

19. R. H. Thaler and C. R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness.* Yale University Press, New Haven, CT, USA, 2008.
20. J. Turland, L. Coventry, D. Jeske, P. Briggs, C. Laing, I. Yevseyeva, and A. van Moorsel. Nudging towards security: Developing an application for wireless network selection for android phones. In *Mobile HCI 2014*. Springer, 2014. submitted.

A Proofs

Proposition 1 *The optimal policy for an influencer a_j is defined by, for any θ_j:*

$$\pi_j^*(\theta_j) = \arg\min_{\eta \in \mathcal{A}_j} \left[\sum_{e \in \mathcal{E}} p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \eta) \right]$$

Proof. Let us show that $\Delta_i^{\pi^*}(e) \leq \Delta_i^{\pi}(e)$, for any environment e and any policy π. By definition of π_j^*, we have, for any θ_j:

$$\begin{aligned} &\sum_{e \in \mathcal{E}} p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j^*(\theta_j)) \\ &\leq \sum_{e \in \mathcal{E}} p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j(\theta_j)) \end{aligned}$$

We can safely multiply by $p(\theta_j)$ on both sides since it is always positive:

$$\begin{aligned} &p(\theta_j) \sum_{e \in \mathcal{E}} p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j^*(\theta_j)) \\ &\leq p(\theta_j) \sum_{e \in \mathcal{E}} p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j(\theta_j)) \end{aligned}$$

Summing over all possible θ_j and factoring, we have:

$$\begin{aligned} &\sum_{\theta_j \in \Theta_j} \sum_{e \in \mathcal{E}} p(\theta_j) p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j^*(\theta_j)) \\ &\leq \sum_{\theta_j \in \Theta_j} \sum_{e \in \mathcal{E}} p(\theta_j) p(e \mid \theta_j) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j(\theta_j)) \end{aligned}$$

Using Bayes theorem, it follows:

$$\begin{aligned} &\sum_{\theta_j \in \Theta_j} \sum_{e \in \mathcal{E}} p(e) p(\theta_j \mid e) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j^*(\theta_j)) \\ &\leq \sum_{\theta_j \in \Theta_j} \sum_{e \in \mathcal{E}} p(e) p(\theta_j \mid e) \sum_{\theta_i \in \Theta_i} p(\theta_i \mid e)\, \delta_i(\theta_i, e, \pi_j(\theta_j)) \end{aligned}$$

Using Definition 2 we conclude after refactoring that:

$$\sum_{e \in \mathcal{E}} p(e) \Delta_i^{\pi^*}(e) \leq \sum_{e \in \mathcal{E}} p(e) \Delta_i^{\pi}(e)$$

B Tables

Table 2. Impact of all effects for any context c and the observation θ_1 in the environments e_1 and e_2, where $\rho_i(e_1, s) = 0$, $\rho_i(e_1, f) = 1$, $\rho_i(e_2, s) = 1$ and $\rho_i(e_2, f) = 0$.

η_{xy}	$v(\theta_1, \eta_{xy}(c), s)$	$v(\theta_1, \eta_{xy}(c), f)$	$\pi_i(\theta_1, \eta_{xy}(c))$	$\delta_i(\theta_1, e_1, \eta_{xy})$	$\delta_i(\theta_1, e_1, \eta_{xy})$
η_{NN}	0.5	0.4	s	0	1
η_{NR}	0.5	0.3	s	0	1
η_{NG}	0.5	0.5	s	0	1
η_{GN}	0.6	0.5	s	0	1
η_{RN}	0.4	0.4	s	0	1
η_{GR}	0.6	0.3	s	0	1
η_{RG}	0.4	0.5	f	1	0
η_{RR}	0.4	0.3	s	0	1
η_{GG}	0.6	0.5	s	0	1

Table 3. Impact of all effects for any context c and the observation θ_2 in the environment e_2, where $\rho_i(e_2, s) = 1$ and $\rho_i(e_2, f) = 0$.

η_{xy}	$v(\theta_2, \eta_{xy}(c), s)$	$v(\theta_2, \eta_{xy}(c), f)$	$\pi_i(\theta_2, \eta_{xy}(c))$	$\delta_i(\theta_2, e_2, \eta_{xy})$
η_{NN}	0.5	0.65	f	0
η_{NR}	0.5	0.55	f	0
η_{NG}	0.5	0.75	f	0
η_{GN}	0.6	0.65	f	0
η_{RN}	0.4	0.65	f	0
η_{GR}	0.6	0.55	s	1
η_{RG}	0.4	0.75	f	0
η_{RR}	0.4	0.55	f	0
η_{GG}	0.6	0.75	f	0