Close the Intention-Behavior Gap via Attitudes: Case Study of the Volitional Adoption of a Two-Factor Authentication Service

Tom Mattson
University of Richmond
tmattson@richmond.edu

Salvatore Aurigemma
University of Tulsa
sal@utulsa.edu

Jie Ren
Fordham University
jren11@fordham.edu

Abstract

Most of the theories used in the behavioral security literature explain the variance in intentions to act securely. Yet, individuals often fail to act on their intentions. This disconnect is referred to as the intention-behavior gap. Most theories propose a single structural path between intentions and actual behaviors with the expectation that individuals will act on their intentions. The purpose of our paper is to investigate this intention-behavior gap in the context of the volitional adoption of information security technologies. To do so, we conducted a two-phased qualitative study of the adoption of a two-factor authentication (2FA) service. In our bottom-up investigation, we discovered emergent themes related to the four functional areas of attitudes (i.e., functional attitude theory). Our paper contributes to the behavioral security literature by suggesting that individuals must change their negative attitudes related to different functional areas to start to reduce the intention-behavior gap.

Keywords: intention-behavior gap, behavioral information security, 2FA services, and functional theory of attitudes

1. Introduction

Behavioral intentions are individuals' desire or their perceived likelihood that they will engage in a specific action (Ajzen, 1991). When asked, individuals typically claim to have high intentions to take precautionary actions to protect themselves from cyber threats but they often fail to act on their intentions (Aurigemma & Mattson, 2019; Liang & Xue, 2009). The disconnect between behavioral intentions and actual behaviors is referred to as the intention-behavior gap (Sheeran, 2002), which may be quite high for information security actions (Crossler, Long, Loraas, & Trinkle, 2014; Jenkins, Durcikova, & Nunamaker Jr., 2021). This gap is problematic in the context of information security because individuals are not safe when they only intend to act. They must convert those intentions into actual actions in order to be protected from cyber threats. By themselves, high intentions offer zero protection against cyber threats.

In the behavioral information security literature, however, scholars most often investigate the antecedents of behavioral intentions to act securely or follow their organizations' policies and procedures (D'Arcy & Herath, 2011; Menard, Bott, & Crossler, 2017; Moody, Siponen, & Pahnila, 2018). They focus on behavioral intentions for several reasons. First, the theories and models used in the literature (e.g., deterrence theory, rational choice theory, protection motivation theory, fear-appeals model, health belief model, technology acceptance model, and the theory of planned behavior) specifically theorize about the constructs (and relationships thereof) that impact behavioral intentions or protection motivation intentions. Second, it is significantly easier to measure behavioral intentions instead of actual behaviors. Scholars cannot use Likert style questions to measure actual adoption rates like they can for adoption intentions. Self-reported actual adoption measures have minimal scientific value. Third, there is no elegant theory that explains the gap between intentions and actual behaviors. It is the sophisticated theoretical explanations that make manuscripts publishable.

As a result, our research community knows very little about the factors that moderate or mediate (partially or fully) the path between behavioral intentions and actual adoption (Anwar, He, Ash, Yuan, Li, & Xu, 2017; Crossler et al., 2014; Jenkins et al., 2021). In his cross-disciplinary meta-analysis, Sheeran (2002) found that only a small fraction (28%) of the explained variance in actual behaviors was explained by their behavioral intentions. If that result is generalizable to information security behaviors, then that means most of the variance in actual security behaviors is not explained by the single structural path between intentions and actual behaviors that are common across most of the theories used in the behavioral security literature.

URI: https://hdl.handle.net/10125/103135
978-0-9981331-6-4

Therefore, behavioral security scholars might be focused on constructing and testing behavioral models and theories that explain the variance associated with the smaller part of the problem (antecedents of intentions) of convincing individuals (personal users or employees) to act securely. The larger part of the problem might be translating those high intentions into actual security actions. This issue is analogous to convincing an individual to "intend" to start exercising. That might be easier than actually convincing them to follow through on their intentions such that they actually start exercising regularly. Similarly, convincing individuals that they need to take precautionary actions might be easier than convincing them to actually practice safe computing practices regularly. It is difficult to be cautious, mindful, and diligent on a daily basis even though they may have high intentions to do so. As such, we address the following research question in our paper:

RQ: What factors inhibit or facilitate individuals from acting on their behavioral intentions to adopt security technologies volitionally?

To answer this research question, we performed a qualitative study of individuals' adoption of a 2FA service. In our qualitative analyses, we found themes related to attitudes as explanations for not following through on their intentions. More specifically, we found that the functional attitude theory (FAT) helped explicate why individuals did not follow through on their intentions (Katz, 1960). Our participants' negative attitudes resulted from utilitarian (i.e., low priority to security actions), knowledge (i.e., lack of knowledge or dismissive of new information), value-expressive (i.e., secure computing is something that is not valued), and ego defensive (i.e., rationalizing their non-action) functional mechanisms. Therefore, our results suggest that understanding attitudes and their different functional components help explain why certain individuals do not follow through on their behavioral intentions.

2. Literature Review

In this section, we review relevant and selected literature on the intention-behavior gap along with the literature on attitudes. The gap is related to all types of actions (not just security or technology actions) and so is the attitude construct. As a result, we review a broad range of literature in this section.

2.1. Intention-Behavior Gap

Many behavioral theories were specifically developed to explain the variability in behavioral intentions. The general idea is that individuals first form their intentions and then they act on those intentions (Ajzen & Kruglanski, 2019). Therefore, it is important to understand how individuals form their intentions to act. These behavioral theories propose direct, moderating, and/or mediating (partial or full) relationships between their core constructs and behavioral intentions as the primary outcome variable. These models and theories then propose a single linear relationship between behavioral intentions and actual behaviors (Crossler et al., 2014; Jenkins et al., 2021). This single path between intentions and actual behaviors stems from the theory of planned behavior and its predecessor the theory of reasoned action, which suggests a rational process where individuals first intend to act and then actually perform the intended action or set of actions (Ajzen, 1991).

Often, however, there is a gap between behavioral intentions and actual behaviors, which has been reported across many different actions (Sheeran, 2002). For instance, we see an intention-behavior gap in exercising, ethical consumerism, and green consumption (Carrington, Neville, & Whitwell, 2014; Nguyen, Nguyen, & Hoang, 2019; Papies, 2017; Sheeran & Webb, 2016). In an information security context, we see evidence that suggests personal users and employees intend to practice safe computing practices (i.e., individuals have high intentions to act securely) but they still fail to act (as evident by the low adoption rates of many security-related technologies and recommended security best practices) (Anwar et al., 2017; Crossler et al., 2014; Jenkins et al., 2021; van Bavel, Rodriguez-Priego, Vila, & Briggs, 2019).

Sheeran (2002) reports that behavioral intentions only account for 28% of the explained variance in actual behaviors in their meta-analysis of studies across a variety of actions. Obviously, any variance explained metric is non-linearly related to the magnitude of the effect sizes (Hunter & Schmidt, 1990, p. 199). Therefore, explaining 28% of the variance is not necessarily an indication of a low effect size for the path coefficient between intentions and actual behaviors. However, this statistic does highlight (at least on the surface) that more research is needed to further our understanding of the intention-behavior gap. Furthermore, given the difficulties that individuals have in converting their intentions into actual behaviors across many different contexts (Sheeran & Webb, 2016; Wanberg, Zhu, & Van Hooft,

2010), it seems unlikely that a single (non-moderated and non-mediated) structural path between intentions and actual behaviors captures the behavioral complexity associated with actually performing security behaviors.

From the cross-disciplinary literature on the intention-behavior gap, we know several factors associated with this gap. First, individual-level characteristics impact the conversion of intentions into actions (Sheeran, 2002; Sheeran & Webb, 2016). For instance, individuals' personality types along with cognitive perceptions concerning many situational and environmental factors impact how likely they are to follow through on their behavioral intentions (Margolis & Lyubomirsky, 2020; Pfeffer, Englert, & Mueller-Alcazar, 2020). Interestingly, the prior literature reports that high self-efficacy does not always translate into higher actual actions (Sheeran, 2002; Sheeran & Webb, 2016) even though self-efficacy has been demonstrated to be strongly related to the formation of behavioral intentions in a variety of contexts (Bandura, 1986).

Second, the type of behavior that individuals intend to perform also impacts the intention-behavior gap (Sheeran, 2002; Sheeran & Webb, 2016). Certain behavioral intentions are related to single actions (e.g., "I intend to take the bus tomorrow"), whereas others require multiple actions (e.g., "I intend to reduce my carbon footprint this month"). Single action behavioral intentions tend to have a smaller intention-behavior gap relative to multiple action behavioral intentions (Sheeran, 2002). Many information security behaviors such as periodically changing passwords, regularly updating anti-malware software, and continuously patching operating systems require multiple actions, which could result in a rather high intention-behavior gap in this context.

Third, the type of intention (and properties thereof) has also been reported to help explicate the intention-behavior gap (Gollwitzer, 1999; Gollwitzer & Sheeran, 2006; Sheeran, 2002). Generic intentions (e.g., "I intend to be environmentally responsible") are less effective at bridging the intention-behavior gap relative to specific implementation intentions (e.g., "I intend to be environmentally responsible by reusing containers at the grocery store tomorrow") (Gollwitzer & Sheeran, 2006). Specific implementation intentions outline the when, where, and how that an individual will act on their intention, which makes those types of intentions more effective than generic intentions in terms of acting on those intentions (Gollwitzer, 1999; Gollwitzer & Sheeran, 2006).

In the behavioral information security literature, Anwar et al. (2017) argue and demonstrate empirically that individual gender differences impact individuals' actual security-related actions. Crossler et al. (2014) found that perceived costs (effort) associated with the security action helped explain the intention-behavior gap with an organization's bring your own device policy. Building on Crossler et al. (2014)'s results, Jenkins et al. (2021) found that effort associated with security actions moderated the path between intentions and actual behaviors. They argued that individuals have conflicting goals related to effort (both minimizing and maximizing) that impacted the intention-behavior gap for security actions (Jenkins et al., 2021).

Many security technologies require effort to install, configure, and use on an on-going basis (Aurigemma & Mattson, 2019; Warkentin, Johnston, Shropshire, & Barnett, 2016). This required effort, however, is often not fully known when individuals form their initial behavioral intentions and motivation intentions to adopt a security technology or perform a security-related action. That is one of the reasons why individuals' intentions and motivations are only the starting point for actual actions in this context. An individual's attitudes towards an action drives effort to perform the action (van Schie, Martijn, & Van Der Pligt, 1994), which may not be fully captured in their behavioral intentions. We proffer that having a negative or destructive attitude will result in less effort to act on their behavioral intentions. For instance, individuals who have a negative attitude concerning the benefits of patching their operating system will probably devote less effort towards patching their operating system relative to individuals who have a more positive attitude towards the action.

2.2. Attitudes

Attitudes refer to an internal evaluation of a system, action, or idea (Ajzen, 1991; Petty, Priester, & Wegener, 2014). These evaluations typically range from favorable to unfavorable or positive to negative (Fishbein & Ajzen, 1977). As a result, attitudes impact a variety of behavioral intentions including technology adoption intentions and security action intentions (Blut, Chong, Tsiga, & Venkatesh, 2021; Dwivedi, Rana, Jeyaraj, Clement, & Williams, 2019; Herath & Rao, 2009). The general idea is that negative (unfavorable) attitudes will lead to low behavioral intentions and positive (favorable) attitudes will lead to high behavioral intentions (Ajzen, 1991; Ajzen & Kruglanski, 2019). That is, it is difficult to convince individuals to form high behavioral intentions if they

have a negative (unfavorable) attitude towards the behavior.

The functional attitude theory (FAT) suggests that attitudes serve particular psychological functions for individuals (Katz, 1960). This theory posits that to understand an individual's attitudes, we must understand the root cause (functional area) behind their attitude. In general, the FAT proposes that there are four functional areas of attitudes: 1) utilitarian whereby attitudes are formed based on rewards and punishments, 2) knowledge whereby attitudes are formed based on information (or lack thereof), 3) value-expressive whereby attitudes are formed based on individual values and concepts of the self, and 4) ego defensive whereby attitudes are formed based on protection mechanisms due to external threats. In order to change an individual's attitude, the FAT argues that psychologists, managers, academics, and consultants must focus on a specific functional area (Poels & Dewitte, 2019). For instance, appealing to the ego-defensive function might be used to influence individuals who practice unsafe computing practices but perceive themselves as practicing safe computing. To change this individual's attitude towards security actions, the FAT would suggest appealing to their self-beliefs that they believe that they are practicing safe computing (i.e., go from their perceived beliefs that they are practicing safe computing practices to safer even though they are practicing unsafe computing practices). That is, we have to match the marketing message with the individual's functional attitudinal dimension to maximize its effectiveness.

The attitude to behavioral intention path has been consistently reported (positive effect) in the prior information systems literature with a variety of technology-related actions (Bélanger, Collignon, Enget, & Negangard, 2017; Blut et al., 2021; Dwivedi et al., 2019; Herath & Rao, 2009). Our qualitative analyses also suggest that attitudes help explain the intention-behavior gap in this security context for a couple of reasons. First, the attitude that gets formed when individuals establish their initial behavioral intentions may differ from the attitude that is needed to convert intentions into actual actions. For instance, an individual may have a positive attitude towards adopting an anti-malware application when they form their initial adoption intentions. However, as they investigate what is required to actually adopt or use the software regularly post-adoption, their utilitarian or ego defensive attitudes may change from mostly positive to mostly negative, which might explain why they do not follow through on their initially high adoption intentions.

Second, Jenkins et al. (2021) and Crossler et al. (2014) proposed that perceived effort to perform the action was a significant moderator of the intention to actual behavior path. We suggest that effort contributes to the formation of positive or negative functional attitudes towards the actual behavior. The more difficult the action is to perform; the more negative or unfavorable an individual's attitudes will be towards the action. The more unfavorable the attitude, the less likely they will be to invest their time and energy to perform the precautionary action that requires effort to perform.

3. Research Design and Methods

To investigate the intention-behavior gap for security actions, we conducted a qualitative study of the volitional adoption (or lack thereof) of 2FA services. The 2FA service we investigated was linked to our participants' University email accounts. This particular 2FA service was relatively easy to setup (i.e., just a selection in their user profiles that required only a few mouse clicks to set-up). However, it does require a moderate amount of effort to use on an on-going basis (post-adoption use) because each email log-in attempt required entering a code that was sent via text message to their cell phones.

In our study, we focused on personal users instead of employees in organizations because personal users do not have a set of information security policies (ISPs) that mandate them to convert their intentions into actions, which takes some (not all) of the individual agency and autonomy out of their security-related decisions. Personal users are not subjected to these types of organizational constraints (Kam, Mattson, & Goel, 2020; Liang, Xue, Pinsonneault, & Wu, 2019). Therefore, the decision-making autonomy associated with personal users makes them a great group of users to investigate the intention-behavior gap based on their own volitional choices to act or to not act on their intentions to do so.

We followed a two-phase research design. The first phase provided a description of the authentication problem (dangers associated with a compromised account) and a proposed solution (2FA service) via a video. The proposed solution for our study was a 2FA service linked to our research participants' University email accounts. This particular 2FA service was implemented at their University several months before the start of our study, but none of our participants had taken the time to configure their email accounts with this 2FA service prior to participating in our study.

After we provided the participants with a description of the problem and the proposed solution, we measured adoption intentions and other constructs such as self-efficacy and attitudes as well as other demographic information. We adapted our measures from Johnston & Warkentin (2010). We measured adoption intentions because we had to distinguish between research participants who had differing levels of adoption intentions to determine how wide the intention-behavior gap was for our research participants. Next, we gave our research participants one week to actually adopt the security technologies.

After one week, we conducted the second phase of our study. Here, we objectively captured whether each participant adopted the technology or not. To do this objectively as opposed to subjectively via self-reported Likert items, we checked with the technology department at the University to determine if they had configured the 2FA service. This objective measure removed many of the problems associated with research participants lying about their actual adoption of the technologies. After we objectively determined actual adoption, we anonymized the data for analysis. Finally, based on whether they actually adopted the security technology or not, we electronically asked them an open-ended question (i.e., "if they adopted it, why" or "if they did not adopt it, why not?"). The responses ranged from short phrases to a few sentences. Those free-form responses were the primary data used in our study.

Our sample consisted of business school students from a private University in the Midwest portion of United States. Many of the complaints about using students in academic research are associated with scholars attempting to generalize their findings from students to employees or managers in organizations (Bello, Leung, Radebaugh, Tung, & Van Witteloostuijn, 2009). For our study related to individual personal users, students are not subjected to any organizational level policies that might spillover into their decision-making in their personal computing environments, which makes them an acceptable sample for our study. We recruited 382 students and 343 completed both phases of our 2FA study. We compared demographic differences (e.g., gender, grade point average, and major within the business school) between those participants who completed both phases and those who did not. We did not notice any demographic differences. Table 1 contains the demographic information for our sample of participants who completed both phases of the study.

Table 1. Demographics

	2FA Service
Total Sample Size	343
Actual Adopters	121
Participants by Age	
18-20	185
21-24	144
>25	14
Participants by Gender	
Female	178
Male	165
Participants by Grades	
<3.0	96
3-3.5	127
>3.5	120
General Computer Knowledge (7-point)	3.133

4. Results

Qualitative research is not based on a single analytical approach because many different analytical methods have emerged over time (Flick, 2009, p. 306). However, one common component of most qualitative studies is coding free-form responses from research participants. A code is typically a word or short phrase that captures the essence of language-based data (Eisenhardt, 1989; Vaast et al., 2013). During analysis, multiple levels of codes are typically developed in a non-linear, iterative manner. Each coding level helps uncover patterns and potential relationships. In our paper, we followed an iterative three-level coding process consistent with Eisenhardt (1989) and Vaast et al. (2013). Our process included: 1) open coding (not grounded in any theory or set of constructs), 2) axial (informed by the literature but still open to data-driven emergent themes), and 3) selective coding (narrow set of codes consistent with one or more specific theories). Our process initially followed a grounded approach where our open level of coding was not informed by any pre-existing theory (Strauss, 1988). However, our axial and selective coding processes narrowed our open codes by using the prior literature to inform our analyses.

We used multiple coders to code our data consistent with the prior research (Mattson, 2017; Vaast et al., 2013). We first coded 50 free-form text responses together to develop a consistent process among the different coders. After the process was refined, the three coders then coded 50 different data points individually to determine inter-rater reliability. All discrepancies with those observations were discussed and resolved collectively. The inter-rater reliabilities between each pair of coders were 0.91, 0.82, and 0.86 respectively, which indicates that our process was consistent across all three coders.

We started with a series of open coding rounds, which resulted in many themes related to individual differences, the threat, and the 2FA technology. Individuals identified that they were not terribly concerned about the threat of having their email accounts compromised. Others mentioned that they would take precautionary actions such as configuring the 2FA service only after their accounts were compromised. This counter-productive strategy was quite prevalent in our sample of personal users. Until the threat personally impacted them, they indicated negative attitudes towards taking any precautionary action. Other participants questioned whether the 2FA service would actually be effective at protecting their email accounts (similar to response efficacy). A few of the other most prevalent open codes were the following: not a priority, not caring about threats, experiential bias, delaying and procrastination, and not convenient. Interestingly, many of the participants who indicated an indifferent attitude or an experiential bias still indicated a relatively high intention to adopt the 2FA service.

Many individuals who did not adopt the 2FA service mentioned that they had "no time" to adopt it. The "no-time" excuse is troubling because configuring the 2FA service probably takes only one- or two-minutes with a handful of mouse clicks to setup and configure. Contrarily, the research participants who adopted the 2FA service mentioned that only minimal time and effort was needed to adopt it. The non-adopters felt that they were better off spending those few minutes doing something else. The "no-time" excuse might result in forming a negative or an unfavorable attitude towards actually performing the security action.

After our open coding, we consulted the literature to perform a series of axial coding rounds. These rounds of coding were informed by our open codes, the free-form responses (the data), and the prior literature. Here, we identified constructs related to the threat (vulnerability, severity, and apathy) along with, indifference, defensive mechanisms, limited personal resources, biases, low technology response efficacy, and solution hubris. Our analysis of these axial codes revealed that many of them were indirectly or directly related to the formation of an individual's positive or negative attitudes towards the 2FA service. The individuals who adopted the 2FA service had positive opinions and attitudes towards the technology while the non-adopters had negative opinions and attitudes. However, the attitudes (positive or negative) were not all related to the same functional component of the 2FA service and the identity management threat. As a result, we narrowed our focus down to the FAT, which informed our selective coding efforts.

Table 2 contains our selective codes and definitions related to the FAT. Different participants focused on different functional areas in their free-form responses, which impacted their attitudes towards the 2FA service. For instance, some responses were utilitarian (e.g., "2FA not needed" or "2FA protects my email") and others were based on the knowledge functional area (e.g., "I had no idea my account was at risk"). Particularly for the non-adopters, many developed their negative or unfavorable attitudes towards the 2FA service based on rationalizing their non-action (ego defensive). Other participants valued the ease of access of their email more than the added security that came with the 2FA service (negative value-expressive attitude). For many of the adopters, however, they had the opposite value proposition whereby they valued security over convenience when they formed their positive attitude regarding the 2FA service.

Table 2. Selective Codes and Definitions of Attitudes

Dimensions	Definition
Utilitarian	The utilitarian attitude provides general approach or avoidance tendencies
Knowledge	The knowledge attitude organizes and interprets new information
Ego-defensive	The ego-defensive attitude protects self-esteem
Value-expressive	The value-expressive attitude expresses central values or beliefs
Source: Katz (1960)	

We conducted proportion tests for the counts of each functional attitude between adopters and non-adopters. Table 3 shows these results. All functional areas were significant except for the knowledge

function. Therefore, having more information did not contribute to the intention-behavior gap in our data. However, many of the adopters found utility in the 2FA service, which contributed to their positive attitudes and helped them convert their intentions to actual adoption. The non-adopters found significantly more reasons to justify their inaction and their negative attitudes via the ego-defensive function of attitudes.

Table 3. Counts by Adopters versus Non-adopters

Dimensions	Adopters	Non-adopters	P-value
Utilitarian	85	10	***
Ego-defensive	4	137	***
Value-expressive	12	42	*
Knowledge	20	33	NS
Totals	121	222	

*** p<0.001, * p<0.05, NS not significant

We next compared the coded FAT dimensions for those who indicated high versus low behavioral intentions to adopt the 2FA service. We used below and above the mean scores on the behavioral intentions measures to determine low intentioned and high intentioned individuals. We then conducted proportion tests to compare the counts between the two intention groups. Tables 4 & 5 show these results for the adopters and the non-adopters respectively. For the adopters, the proportion of high intentioned individuals had a positive utilitarian attitude was greater than the proportion in the low intentioned group. All of the other coded attitude groups were not significantly different between their intention levels.

Table 4. Counts of Adopters by Intentions

Dimensions	High Intention	Low Intention	P-value
Utilitarian	70	15	*
Ego-defensive	2	2	NS
Value-expressive	7	5	NS
Knowledge	13	7	NS
Total	92	29	

* p<0.05, NS not significant

For the non-adopters, the proportion of low intentioned research subjects who had a negative ego-defensive attitude was greater than the proportion in the high intentioned group. All of the other coded attitude groups were not significantly different between their intention levels for the non-adopters. Interestingly, there were more non-adopters who had high intentions to adopt relative to low-intentions.

Table 5. Counts of Non-Adopters by Intentions

Dimensions	High Intention	Low Intention	P-value
Utilitarian	8	2	NS
Ego-defensive	67	70	*
Value-expressive	29	13	NS
Knowledge	19	14	NS
Total	123	99	

* p<0.05, NS not significant

Our selective FAT codes and our initial counts along with the proportion tests suggest that different functional components of attitudes related to the action moderate the intention-behavior gap. Figure 1 displays a diagram of a potential research model with different moderating relationships.

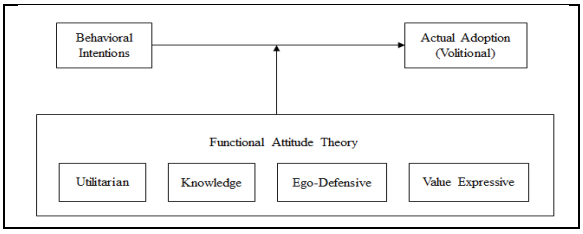

Figure 1. Research Model

To test the moderating effects, we ran a series of logistic regressions. We tested each coded attitude functional area along with its interaction effect with behavioral intentions separately. Table 6 displays those results. Across the four models, we see a consistent significant and positive association between intention of adoption and actual adoption. However, both the value-expressive and knowledge attitudes negatively moderate this positive association (at the 0.1 level). That is, the positive association between intention and actual adoption is weaker when people show either the value-expressive attitude or the knowledge attitude.

Table 6. Logistic Regressions

Actual Adoption Dependent Variable				
	(1)	(2)	(3)	(4)
Intercept	-3.48 ***	-1.91 **	-4.16 ***	-4.23 ***
Intention	0.40 **	0.52 ***	0.80 ***	0.77 ***
Utilitarian	2.06			
Ego-defensive		-2.85		
Value-expressive			1.37	
Knowledge				2.41*
Utilitarian *Intention	0.30			
Ego-defensive *Intention		-0.23		
Value-expressive *Intention			-0.48*	
Knowledge *Intention				-0.48*
AIC	257	257	369	378

* p<0.1, **p<0.05, ***p<0.001

5. Discussion and Conclusion

Secure computing happens only when individuals act on their intentions. That is, individuals who have high intentions to install an anti-malware application on their phone but fail to actually install it results in vulnerable device. Intentions without actions offer zero protection. As such, researchers cannot stop our research at behavioral intentions even if our theoretical explanations generally stop at behavioral intentions. That is only part of the problem (possibly the smaller part of the problem). The prior literature has done a wonderful job applying elegant theories to explain the variability in behavioral intentions but information security scholars have done a worse job explaining the intention-behavior gap (D'Arcy & Herath, 2011; Menard, Bott, & Crossler, 2017; Moody, Siponen, & Pahnila, 2018). Yet, our data indicate that many of our subjects had high intentions to adopt the 2FA service but failed to follow through on those intentions.

Attitudes have been theorized to impact technology adoption intentions across many types of technologies (Bélanger et al., 2017; Blut et al., 2021; Dwivedi et al., 2019; Herath & Rao, 2009). We show in our qualitative data that positive attitudes are also important to convert adoption intentions to actual behaviors. Changing one's attitudes is a powerful way to close the intention-behavior gap. It is difficult to convert intentions into actions if individuals have negative attitudes about the technology across all functional areas. Therefore, we contribute the FAT to the intention-behavior gap literature. This theoretical insight is relevant regardless of the theory that is used to explain the variability in behavioral intentions.

Themes related to negative attitudes (either directly or indirectly) were the most important factors that our research participants identified that contributed to the intention-behavior gap with our 2FA study. Many of our research participants had indifferent or negative attitudes towards our 2FA service. Interestingly, these indifferent or negative attitudes were not apparent when they answered our Likert items regarding attitudes. They were only apparent when we asked them open ended questions regarding why they did not adopt the security software. This pattern suggests that either their attitudes changed as they investigated whether or not to actually adopt the security software or our Likert items did not adequately capture their attitudes towards the 2FA service.

Similarly, some of the qualitative responses as to why they decided not to adopt the 2FA service (e.g., don't like it, don't need it, no desire to install it, already using a different substitute product, and don't see the need for it) made us question the authenticity of the behavioral intention and attitude scores on our 7-point Likert continuums. In their free-from responses, many of our participants indicated that they really had no intention of volitionally adopting the technology now or in the future. However, many of them still responded with a 5 or higher for the behavioral intention and attitude Likert questions. Therefore, social desirability or other issues on these Likert item questions might be problematic for investigating the intention-behavior gap and for investigating the behavioral antecedents of behavioral intentions.

Many of our research participants mentioned that they would install and use the 2FA service only after they were adversely affected by a data breach (i.e., information security is not important until "I" am personally impacted), which is clearly an ineffective mitigation strategy and an unproductive mindset. It is similar to leaving one's car unlocked until it gets

stolen. After their car is stolen then they will think about locking their car doors. As a result, practitioners and academics have to further investigate how to convince individuals that these volitional security actions are important to take before (not after) they are compromised.

A component of the FAT research has focused on the matching hypotheses. The idea of this conjecture is that advertising messages and manipulations for a product or service should match the attitude functions (Herek, 1986; Katz, 1960). A mismatch will result in an unsuccessful or less successful campaign (Teeny, Siev, Briñol, & Petty, 2021). In order to increase the relevance and salience of security campaigns, our results suggest that practitioners might be better off focusing on matching the specific attitude function. Obviously, we did not specifically test this matching hypothesis with our qualitative 2FA study. However, our qualitative results do suggest there were potential mismatches with a few of our open codes. Future research could build off of our results by performing a randomized experiment to specifically look at these potential mismatches.

Like all research, our qualitative 2FA study has a few limitations. First, most of our research subjects were under 25 years old. More research is certainly needed within this age group and outside of this age group to determine how generalizable our findings are. It could be that this younger generation has different functional attitudes regarding security actions and technologies relative to an older demographic but more research is needed to substantiate that claim. We also cannot universally generalize our findings from our convenience sample of younger personal users to other populations of younger personal users.

Second, we investigated personal users as opposed to employees in organizations. Employees in organizations have mandated policies that should be followed, which might impact their attitudes towards performing certain security actions. Employees also have to perform actions that are directly related to their compensation and yearly performance reviews. Those actions will probably result in more positive attitudes along one or more functional areas to convert behavioral intentions to actual actions because those are tied to their paychecks. Following safer computing practices is not typically linked to compensation or performance, which probably impacts their attitudes towards taking the action (even if they "intend" to take the action).

Third, different security technologies require different effort to adopt and use regularly. The 2FA service linked to our study was relatively easy to configure but it did require a fair bit of on-going effort to use on a regular basis. Other technologies might be harder to initially adopt but require less on-going effort to use. Therefore, more research is needed to see if our 2FA context is generalizable to other technologies. Different technologies might result in a different pattern of functional attitudes.

6. References

Ajzen, I. (1991). The Theory of Planned Behavior. Organizational Behavior and Human Decision Processes, *50*(2), 179-211.

Ajzen, I., & Kruglanski, A. W. (2019). Reasoned Action in the Service of Goal Pursuit. Psychological Review, 126(5), 774-786.

Anwar, M., He, W., Ash, I., Yuan, X., Li, L., & Xu, L. (2017). Gender Difference and Employees' Cybersecurity Behaviors. Computers in Human Behavior, 69(4), 437-443.

Aurigemma, S., & Mattson, T. (2019). Generally Speaking, Context Matters: Making the Case for a Change from Universal to Particular ISP Research. Journal of the Association for Information Systems, 20(12), Article 7.

Bandura, A. (1986). The Explanatory and Predictive Scope of Self-Efficacy Theory. Journal of Social and Clinical Psychology, 4(3), 359-373.

Bélanger, F., Collignon, S., Enget, K., & Negangard, E. (2017). Determinants of early conformance with information security policies. Information & Management, 54(7), 887–901.

Bello, D., Leung, K., Radebaugh, L., Tung, R. L., & Van Witteloostuijn, A. (2009). From the editors: Student samples in international business research. Journal of International Business Studies, 40(3), 361-364.

Blut, M., Chong, A., Tsiga, Z., & Venkatesh, V. (2021). Meta-Analysis Of The Unified Theory Of Acceptance And Use Of Technology (UTAUT): Challenging Its Validity And Charting A Research Agenda In The Red Ocean. Journal of the Association for Information Systems, Forthcoming.

Carrington, M. J., Neville, B. A., & Whitwell, G. J. (2014). Lost in Translation: Exploring the Ethical Consumer Intention-Behavior Gap. Journal of Business Research, 67, 2759-2767.

Crossler, R. E., Long, J. H., Loraas, T. M., & Trinkle, B. S. (2014). Understanding Compliance with Bring Your Own Device Policies Utilizing Protection Motivation Theory: Bridging the Intention-Behavior Gap. Journal of Information Systems, 28(1), 209-226.

D'Arcy, J., & Herath, T. (2011). A Review and Analysis of Deterrence Theory in the IS Security Literature: Making Sense of the Disparate Findings. European Journal of Information Systems, 20(6), 643-658.

Dwivedi, Y. K., Rana, N. P., Jeyaraj, A., Clement, M., & Williams, M. D. (2019). Re-examining the unified

theory of acceptance and use of technology (UTAUT): Towards a revised theoretical model. Information Systems Frontiers, 21(3), 719–734.
Eisenhardt, K. M. (1989). Building theories from case study research. Academy of Management Review, 14(4), 532– 550.
Fishbein, M., & Ajzen, I. (1977). Belief, attitude, intention, and behavior: An introduction to theory and research. Philosophy and Rhetoric, 10(2).
Flick, O., (2009). An Introduction to Qualitative Research: Sage Publications.
Gollwitzer, P. M. (1999). Implementation Intentions: Strong Effects of Simple Plans. American Psychologist, 54(7), 493-503.
Gollwitzer, P. M., & Sheeran, P. (2006). Implementation Intentions and Goal Achievement: A Meta-Analysis of Effects and Processes. Advances in Experimental Social Psychology, 38, 69-119.
Herath, T., & Rao, H. R. (2009). Protection motivation and deterrence: a framework for security policy compliance in organisations. European Journal of Information Systems, 18(2), 106–125.
Herek, G. M. (1986). The Instrumentality of Attitudes: Toward a Neofunctional Theory. Journal of Social Issues, 42(2), 99–114.
Hunter, J. E., & Schmidt, F. L. (1990). Methods of Meta-Analysis: Correcting Error and Bias in Research Findings. Newbury Park, CA: Sage.
Jenkins, J. L., Durcikova, A., & Nunamaker Jr., J. F. (2021). Mitigating the Security Intention-Behavior Gap: The Moderating Role of Required Effort on the Intention-Behavior Relationship. Journal of the Association for Information Systems, 22(1), 246-272.
Johnston, A. C., & Warkentin, M. (2010). Fear appeals and information security behaviors: An empirical study. MIS Quarterly, 34(3), 549-566.
Kam, H.-J., Mattson, T., & Goel, S. (2020). A Cross Industry Study of Institutional Pressures on Organizational Effort to Raise Information Security Awareness. Information Systems Frontiers, 22, 1241-1264.
Katz, Daniel (1960). The Functional Approach to the Study of Attitudes. Public Opinion Quarterly, 24(2), 163-204.
Liang, H., & Xue, Y. (2009). Avoidance of Information Technology Threats: A Theoretical Perspective. MIS Quarterly, 33(1), 71-90.
Liang, H., Xue, Y., Pinsonneault, A., & Wu, Y. A. (2019). What Users Do Besides Problem-Focused Coping When Facing IT Security Threats: An Emotion-Focused Coping Perspective. MIS Quarterly, 43(2), 373-394.
Margolis, S., & Lyubomirsky, S. (2020). Experimental Manipulation of Extraverted and Introverted Behavior and Its Effects on Well-Being. Journal of Experimental Psychology: General, 149(4), 719-731.
Mattson, T. (2017). Noise or Quality? Cross-Nested Hierarchical Effects of Culture on Online Ratings. Communication of the Association of Information Systems, 40(1), Article 25.
Menard, P., Bott, G. J., & Crossler, R. E. (2017). User Motivations in Protecting Information Security: Protection Motivation Theory Versus Self-Determination Theory. Journal of Management Information Systems, 34(4), 1203-1230.
Moody, G. D., Siponen, M., & Pahnila, S. (2018). Toward a Unified Model of Security Policy Compliance. MIS Quarterly, 42(1), 285-A22.
Nguyen, H. V., Nguyen, C. H., & Hoang, T. T. B. (2019). Green Consumption: Closing the Intention-Behavior Gap. Sustainable Development, 27(1), 118-129.
Papies, E. K. (2017). Situating interventions to bridge the intention–behaviour gap: A framework for recruiting nonconscious processes for behaviour change. Social and Personality Psychology Compass, 11(7), e12323.
Petty, R. E., Priester, J. R., & Wegener, D. T. (2014). Cognitive processes in attitude change. In Handbook of social cognition (pp. 85–158). Psychology Press.
Pfeffer, I., Englert, C., & Mueller-Alcazar, A. (2020). Perceived Stress and Trait Self-Control Interact with the Intention-Behavior Gap in Physical Activity Behavior. Sport, Exercise, and Performance Psychology, 9(2), 244-260.
Poels, K., & Dewitte, S. (2019). The role of emotions in advertising: A call to action. Journal of Advertising, 48(1), 81-90.
Sheeran, P. (2002). Intention-Behavior Relations: A Conceptual and Empirical Review. European Review of Social Psychology, 12(1), 1-36.
Sheeran, P., & Webb, T. L. (2016). The Intention-Behavior Gap. Social and Personality Psychology Compass, 10(9), 503-518.
Strauss, A. (1998). Qualitative analysis for social scientists. New York: Cambridge University Press.
Teeny, J. D., Siev, J. J., Briñol, P., & Petty, R. E. (2021). A review and conceptual framework for understanding personalized matching effects in persuasion. Journal of Consumer Psychology, 31(2), 382-414.
Vaast, E., Davidson, E., & Mattson, T. (2013). Talking about Technology: The Emergence of a New Actor Category Through New Media. MIS Quarterly, 37(4), pp. 1069-1092.
van Bavel, R., Rodriguez-Priego, N., Vila, J., & Briggs, P. (2019). Using Protection Motivation Theory in the Design of Nudges to Improve Online Security Behavior. International Journal of Human-Computer Studies, 123, 29-39.
van Schie, E. C., Martijn, C., & Van Der Pligt, J. (1994). Evaluative language, cognitive effort and attitude change. European journal of social psychology, 24(6), 707-712.
Wanberg, C. R., Zhu, J., & Van Hooft, E. A. J. (2010). The Job Search Grind: Perceived Progress, Self-Reactions, and Self-Regulation of Search Effort. Academy of Management Journal, 53(4), 788-807.
Warkentin, M., Johnston, A. C., Shropshire, J., & Barnett, W. D. (2016). Continuance of Protective Security Behavior: A Longitudinal Study. Decision Support Systems, 92(1), 25-35.