Poster: Do Users Make Rational Security Decisions?

Elissa M. Redmiles, Michelle L. Mazurek, and John P. Dickerson
Department of Computer Science, University of Maryland

Abstract—**Accurately modeling human decision-making in security is critical to thinking about when, why, and how to recommend that users adopt certain secure behaviors. Here, we present a series of behavioral economics experiments modeling the rationality of end-user security decision-making. We ask participants to make a financially impactful security choice, in the face of transparent risks of account compromise and benefits offered by an optional security behavior (two-factor authentication). We measure the cost and utility of adopting the security behavior via measurements of time spent executing the behavior and estimates of the participant's wage. More than 50% of our participants made rational decisions, and their behavior was boundedly rational: they made decisions based on some risks and context, but not others. Finally, we can model their behavior well (R^2=0.61) as a function of risks, context, and prior behavior.**

Introduction

People's adoption, or rejection, of security behaviors can lead to system-wide consequences.Prior work has proposed two simplified theories of the "human in the loop": a rational actor who chooses to ignore security behaviors because the costs always outweigh the potential losses, and an irrational actor who chooses "dancing pigs over security every time" because they neither understand nor care about security risks [1]. While these simplified models of user behavior can help to provide high-level insights, our aim is to define a more realistic medium between these two extremes: *a semi (or boundedly) rational security actor with predictable and consistent, but not always utility-optimal, behavior based on risks and costs.*

In the work presented here, we seek to provide an empirical, economic examination of the rationality of security behavior in a particular context. We define rationality as utility-optimality: that is, a decision is rational if the utility (gain) from the decision is greater than the costs of enacting the decision. Ultimately, we seek to understand: **How do costs (C), risks (R), and user tendencies and attributes (U) influence: (1) a security decision and (2) whether that decision is rational?** To this end, we construct an experimental system and conduct behavioral-economics experiments (e.g., games) to evaluate and model security decision-making. We find that:

- Users act in accordance with an anchoring effect: they tend to stick with the first security decision they make.
- User decisions to enable 2FA are explained well (pseudo-R^2 =0.612) by their prior behavior, knowledge of costs, and explicit risk judgements and context.
- Users made rational security choices ~50% of the time.
- Users are boundedly rational: they incorporate some knowledge about costs and explicit risks, but not more nebulous risks (e.g., password strength), to inform security decisions.
- In higher-risk conditions, users enable security options more often and make rational decisions more often.
- Users behave more rationally and more securely when protecting assets they already have (endowment effect, moderated by risk).

Consequently, we propose that users can be rationally nudged toward personalized, utility-maximizing security behaviors (or lack of behavior) and that we should allow users more autonomy to make decisions based on transparently communicated risks paired with a push for more data-driven research quantifying those risks. Such solutions will ultimately help end users to make use of their personal, behavioral compliance budgets and maximize market gains from security [2].

Methodology

System. Our experimental system operates like a bank account to which participants have to regularly log in. In each game, participants are assigned to a condition: in one of the endowment conditions participants are given an amount of money and are required to login once every 24 hours to retain their money, while those in one of the earn conditions begin with a small amount of money and have the opportunity to earn more every time they log on. When signing up for a game, participants are offered a security choice: whether to enable two-factor authentication (2FA). Prior to making this choice they are shown explicit risks: risk of being hacked (varies in 8 conditions from 1-50%) and amount of protection (1%, 20% or 50%) offered by adopting 2FA. If they are hacked – probabilistically determined by a script regularly run on the system – participants lose all of the money in their account. At the end of a game, participants receive the real monetary value of the amount left in their account.

Experiments. We recruited 150 workers from Amazon Mechanical Turk (MTurk) to participate in two rounds of an experiment run on our bank system. Each game ran for five days, with a five day break in between, participants could earn up to $5 per round. They spent an average of total of 142 seconds (SD = 35 (s)) logging in to Round 1 (R1) and 158 seconds (SD = 30 (s)) logging into Round 2 (R2).

Variables. Using our system we measure users' 2FA decisions (whether then enabled/didn't enable), their password strength using a data-driven, neural-network meter [3] and their signup and login times (seconds each screen was in focus).

Limitations. Behavioral economics experiments are subject to a number of limitations: participants may behave differently than in real life, our variables for the hack and protect percentages may be unrealistic, and 2FA may not be a representative security behavior. We have done our best to mitigate these limitations by choosing salary and hack percentages close to well-known statistics, and selecting a security behavior that our prior work suggests is a reasonable "middle-of-the-road" in terms of user understanding and adoption [4].

Results

51% of participants in R1 and 56% in R2 chose to enable 2FA. We model R1 decisions to enable 2FA, with logistic regression, as a function of respondent factors (gender, age, education, security behavior intention, internet skill), risks implicitly chosen by the respondent (password strength), and risks and conditions assigned to the participant (H, P, Endow), as well as interactions between these risks and settings.

The model of best fit (pseudo-R^2=0.15) for RD1 shows that those in the endowment conditions are 2.3× more likely to enable 2FA, in line with endowment effects observed in other fields. Those who are shown a higher risk of hacking are more likely to enable 2FA and those who are shown a higher protection from 2FA are also more likely to enable. Further, those in a condition that involved endowment and a higher protection value are even more more likely to enable.

We model R2 decisions as a function of R1 2FA decision, R1 costs (e.g., R1 signup and login times) and the factors above. When modeling R2 behavior as just a function of R1 2FA decision, we find that this model explains 35% of the variance. When modeling R2 behavior as a function of both R1 behavior and R1 costs, we find that we can explain 52% of the variance in R2 behavior. Finally, if we include R2 experimental settings, we explain 61% of behavior variance.

Rationality of Security Decisions. We considered participants to have made a *utility-optimal* decision in the following way: it is utility-optimal to enable 2FA if the cost of doing so is less than the utility that would be gained from 2FA. For those who enabled 2FA in either round, we compute the cost of using 2FA for an individual user as the time it cost them to signup (in hours) plus the sum of the time it cost them to login each time (in hours) times the average U.S. MTurk hourly wage (calculated from recent national survey results), for those who did not enable 2FA we computed costs as 2 times the mean cost of 2FA enablers. We compute the utility of using 2FA as the potential loss (maximum amount they could earn times hack percentage) times the protection gained by using 2FA (P): $U_{2fa} = P[(H) * Max_{bank}]$, where Max_{bank} was $5.

48% of all participants made utility-optimal decisions in R1 and 58% did so in R2. We also find that 64% of those in the medium-risk experiments made a rational choice, none of those in the lowest risk condition, and all of those in the highest risk condition. (2FA was always utility-optimal in the highest risk setting and never in the lowest.) In R2, 69% of 2FA users in the medium-risk experiment make the correct decision, and again all of those in the highest-risk settings and none of those in the lowest-risk settings.

Ultimately, we find that, in R1, 33% of participants in the lowest-risk settings, 48% in the medium risk, and 63% in the highest risk settings make a utility-optimal decision. We observe a learning effect ($\chi^2 = 21.226$, df= 2, p<0.001, V = 0.578 (medium)) with 58% of all participants in the medium-risk experiments making a rational decision in R2, 46% in the lowest-risk settings and 75% in the highest-risk settings. Further, In R1, 61% of those in the endowment condition made utility-optimal choices. We observe no change between the two rounds when comparing utility-optimal decision-making by condition; in R2, 57% of those in the endowment condition.

Finally, we model whether participants made a utility-optimal decision to enable or not enable 2FA, based on the 2× the mean cost of 2FA, for R1 in the same way as we modeled general decision-making above. The model of best fit retains only the hack percentage, setting, and internet skill factors. Those in the endowment setting are 25% more likely to make a utility-optimal choice, while those with higher internet skill are 15% more likely to do so. Further, we see that those who saw a higher hack percentage are more likely to make a *utility-optimal decision* in R1. Finally, we find that these factors explain utility-optimal decision-making in our dataset with a pseudo-R^2 of 0.141. The multi-factor model of best fit for RD2 retains SeBIS, the risk and protection factors, endowment, and interactions between risk and endowment and protection and endowment (psuedo-R^2=0.078); higher SeBIS and hack percent are significantly related to more rational RD2 decisions.

Summary

We find that our participants made utility-optimal decisions more often when faced with higher risks. While perhaps encouraging for corporate high-risk scenarios, this finding also suggests a challenge for day-to-day security, as many of the risks end users confront in daily digital life are less transparent, less monetarily linked, and relatively small. Thus, future work may wish to explore how rationality is affected by different methods of communicating risk, less tangible consequences than the monetary incentives provided in our experiments, and even smaller risks.

Our work supports a nuanced model of the "human-in-the-loop" who is able to some degree to take into account explicit risks and personal costs to make frequently rational decisions, but who struggles to identify less obvious risks (such as those incurred from weak passwords) and relies heavily on prior decisions. This argues for personalized security-behavior recommendations for users tailored based on their costs (e.g., login times), risks (e.g., password strength and other risk factors), and value of their account (e.g., measured through the amount of money stored). Our future work will explore whether such personalized recommendations could provide security benefits and help to avoid large market and personal costs from wasted time and effort on unnecessary behaviors.

References

[1] C. Herley, "So long, and no thanks for the externalities: The rational rejection of security advice by users," in *NPSW*, 2009.

[2] A. Beautement *et al.*, "The compliance budget: Managing security behaviour in organisations," in *NSPW*, 2008.

[3] B. Ur *et al.*, "Design and evaluation of a data-driven password meter," in *CHI*. ACM, 2017.

[4] E. M. Redmiles *et al.*, "How i learned to be secure: a census-representative survey of security advice sources and behavior," in *CCS*. ACM, 2016.

Do Users Make Rational Security Decisions?

Elissa M. Redmiles, Michelle L. Mazurek, and John P. Dickerson
eredmiles@cs.umd.edu

How do costs, risks, and user tendencies influence security decisions and the rationality of those decisions?

Motivation & Method

Users behave insecurely but no mathematical, general behavior model for why.

Users have a limited compliance budget once we have a model we can adjust parameters to help users behave most optimally.

We ran **behavioral economics games on AMT**. Participants could earn up to $5 (mean wage/hr on AMT) by interacting with our system. They played the game (made a decision) up to two times: 125 MTurkers played once and 107 played twice.

Cost is defined as wage-earning time loss

$$C_{2fa} = (T_{signup} + \sum T_{login}) * wage_{mturk}$$

Utility of 2FA is defined the $$$ savings if a hack occurred

$$U_{2fa} = P[(H) * Max_{bank}]$$

Rational behavior achieved when choice utility > cost

Behavioral Economics Experimental System

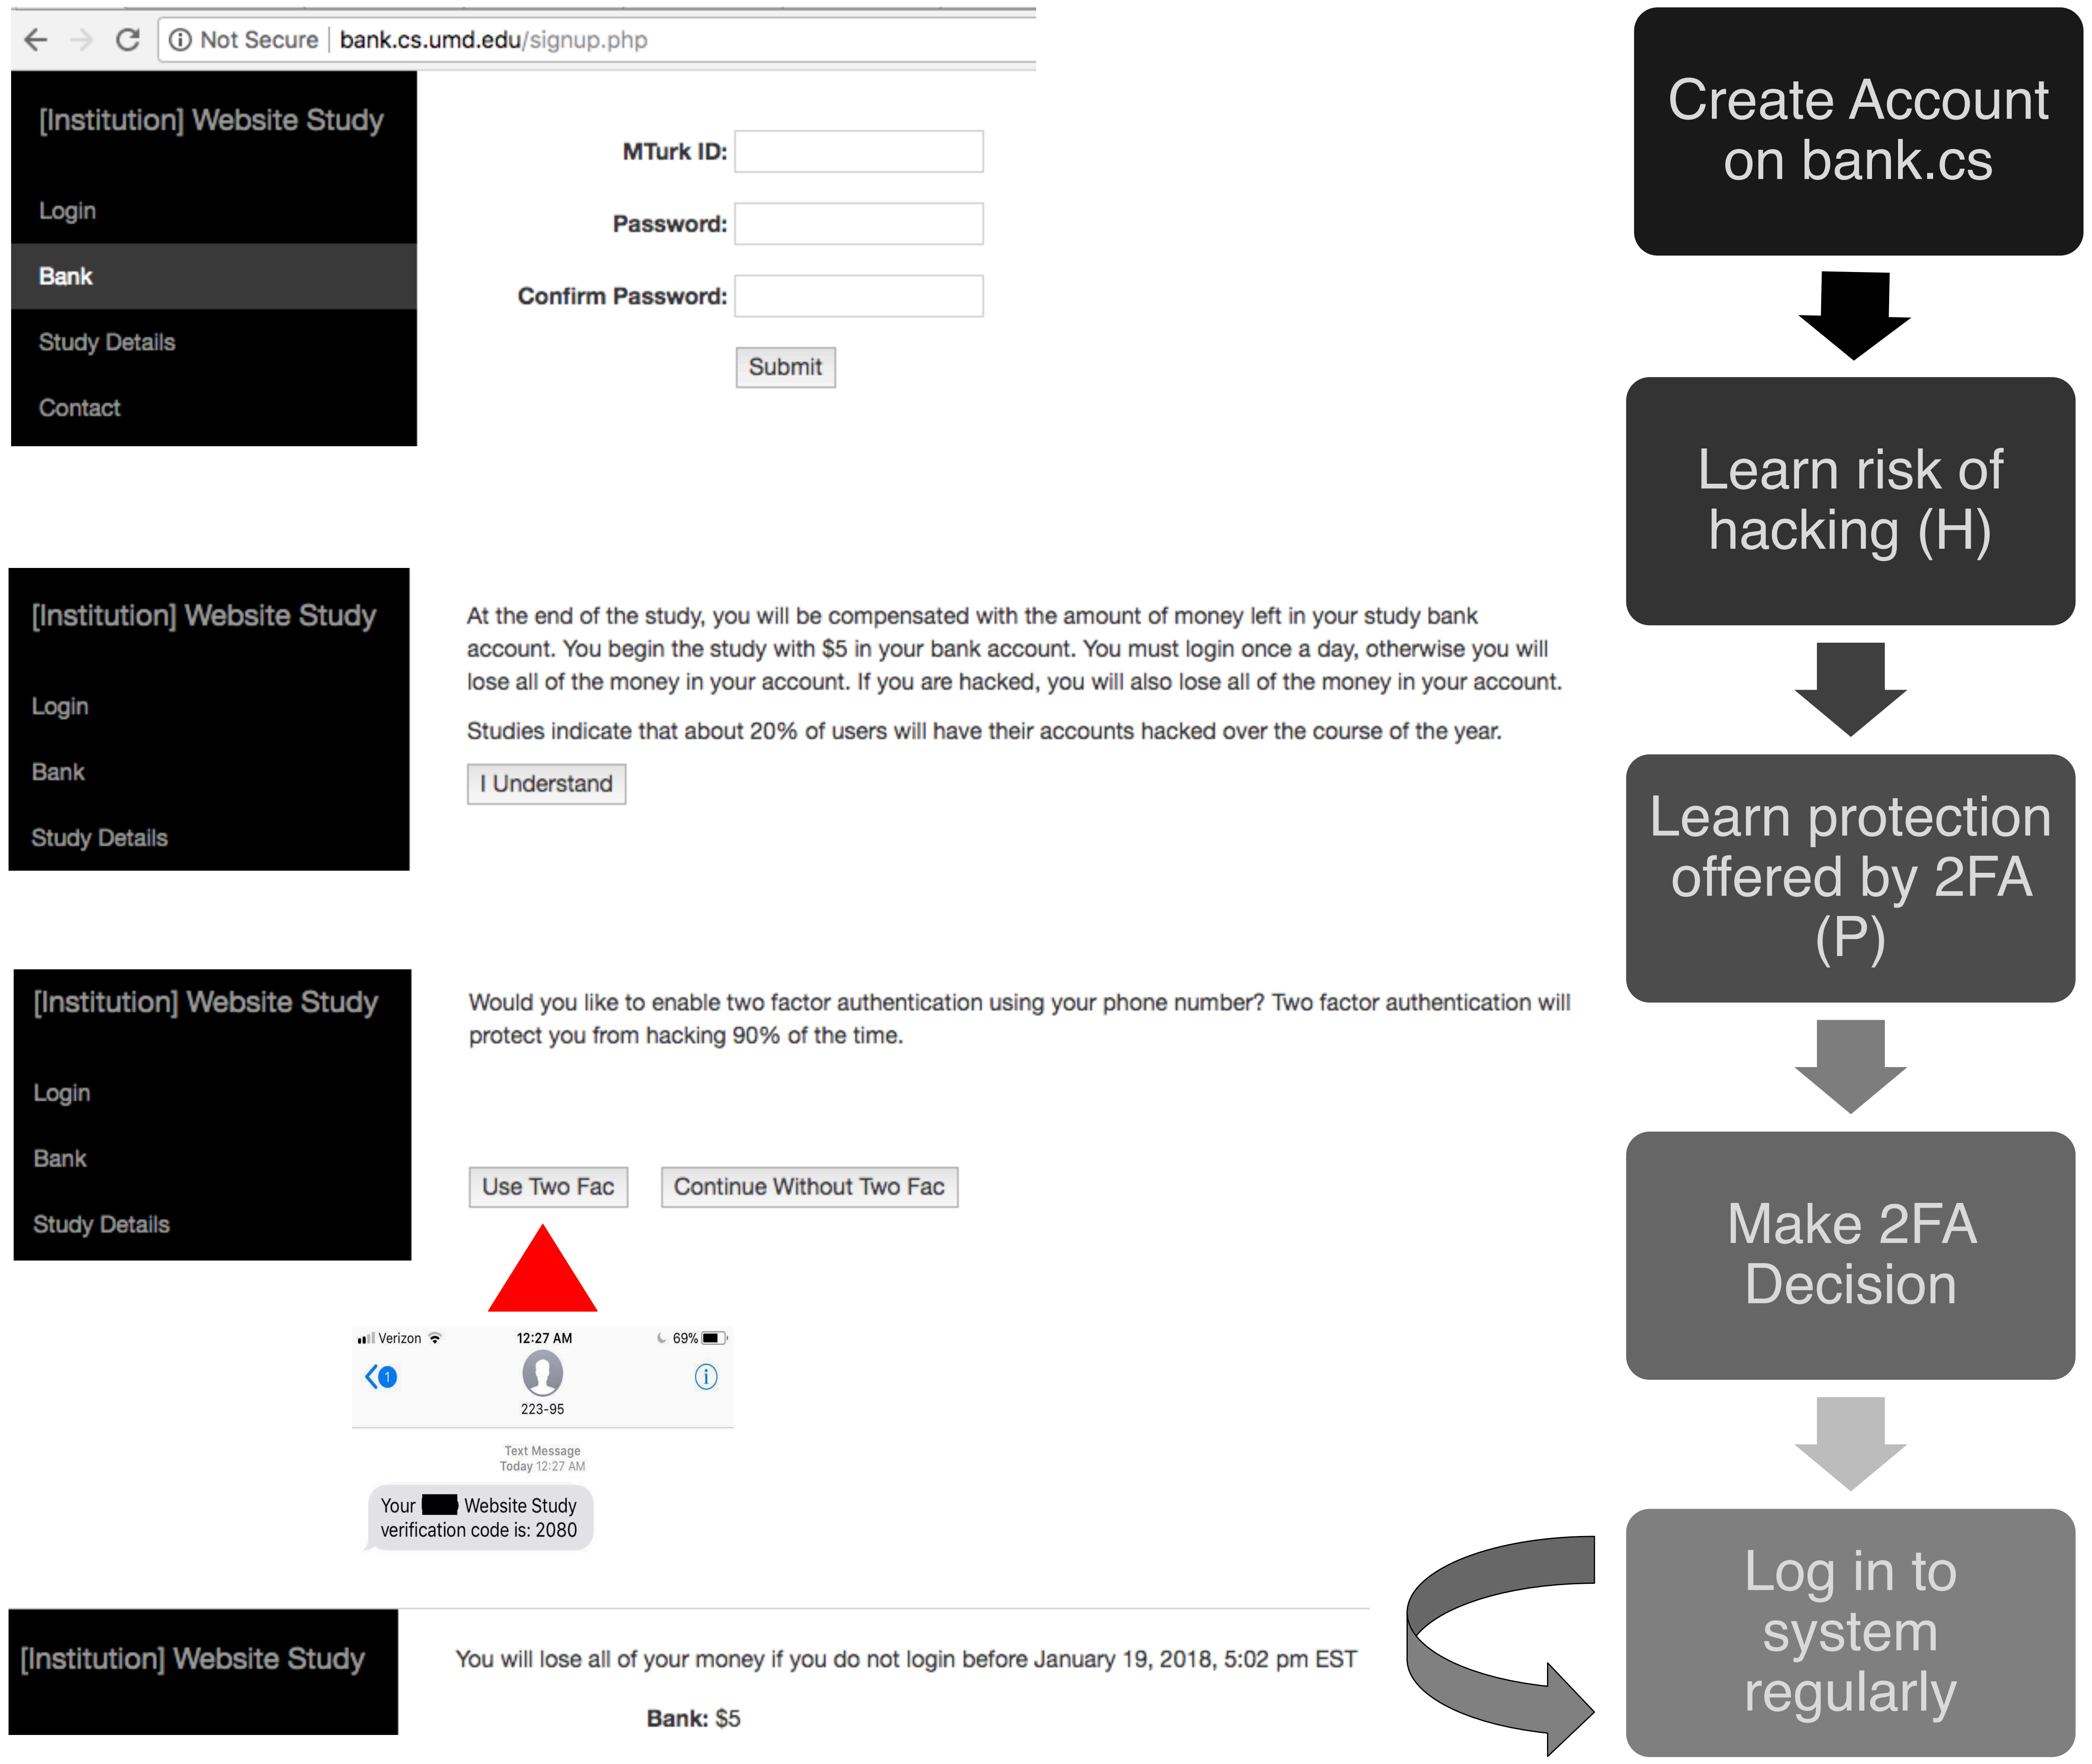

Not Secure | bank.cs.umd.edu/signup.php

[Institution] Website Study

MTurk ID:

Password:

Confirm Password:

[Institution] Website Study

At the end of the study, you will be compensated with the amount of money left in your study bank account. You begin the study with $5 in your bank account. You must login once a day, otherwise you will lose all of the money in your account. If you are hacked, you will also lose all of the money in your account.

Studies indicate that about 20% of users will have their accounts hacked over the course of the year.

[Institution] Website Study

Would you like to enable two factor authentication using your phone number? Two factor authentication will protect you from hacking 90% of the time.

Your [redacted] Website Study verification code is: 2080

[Institution] Website Study

You will lose all of your money if you do not login before January 19, 2018, 5:02 pm EST

Bank: $5

Create Account on bank.cs

Learn risk of hacking (H)

Learn protection offered by 2FA (P)

Make 2FA Decision

Log in to system regularly

Key Findings

Security decisions are explained well (R^2=0.61) by prior behavior, knowledge of costs, endowment effects, explicit risk judgements.

Users made rational decisions ~50% of the time.

Users behavior was **boundedly rational**: they incorporate knowledge about costs and explicit risks, but not more nebulous risks (e.g., password strength) in decisions.

In higher-risk conditions, users enable security options more often and make rational decisions more often.

Users exhibit **endowment effects:** they behave more rationally & more securely when protecting existing assets.

Users act in accordance with an **anchoring effect**: they tend to stick with the first security decision they make.

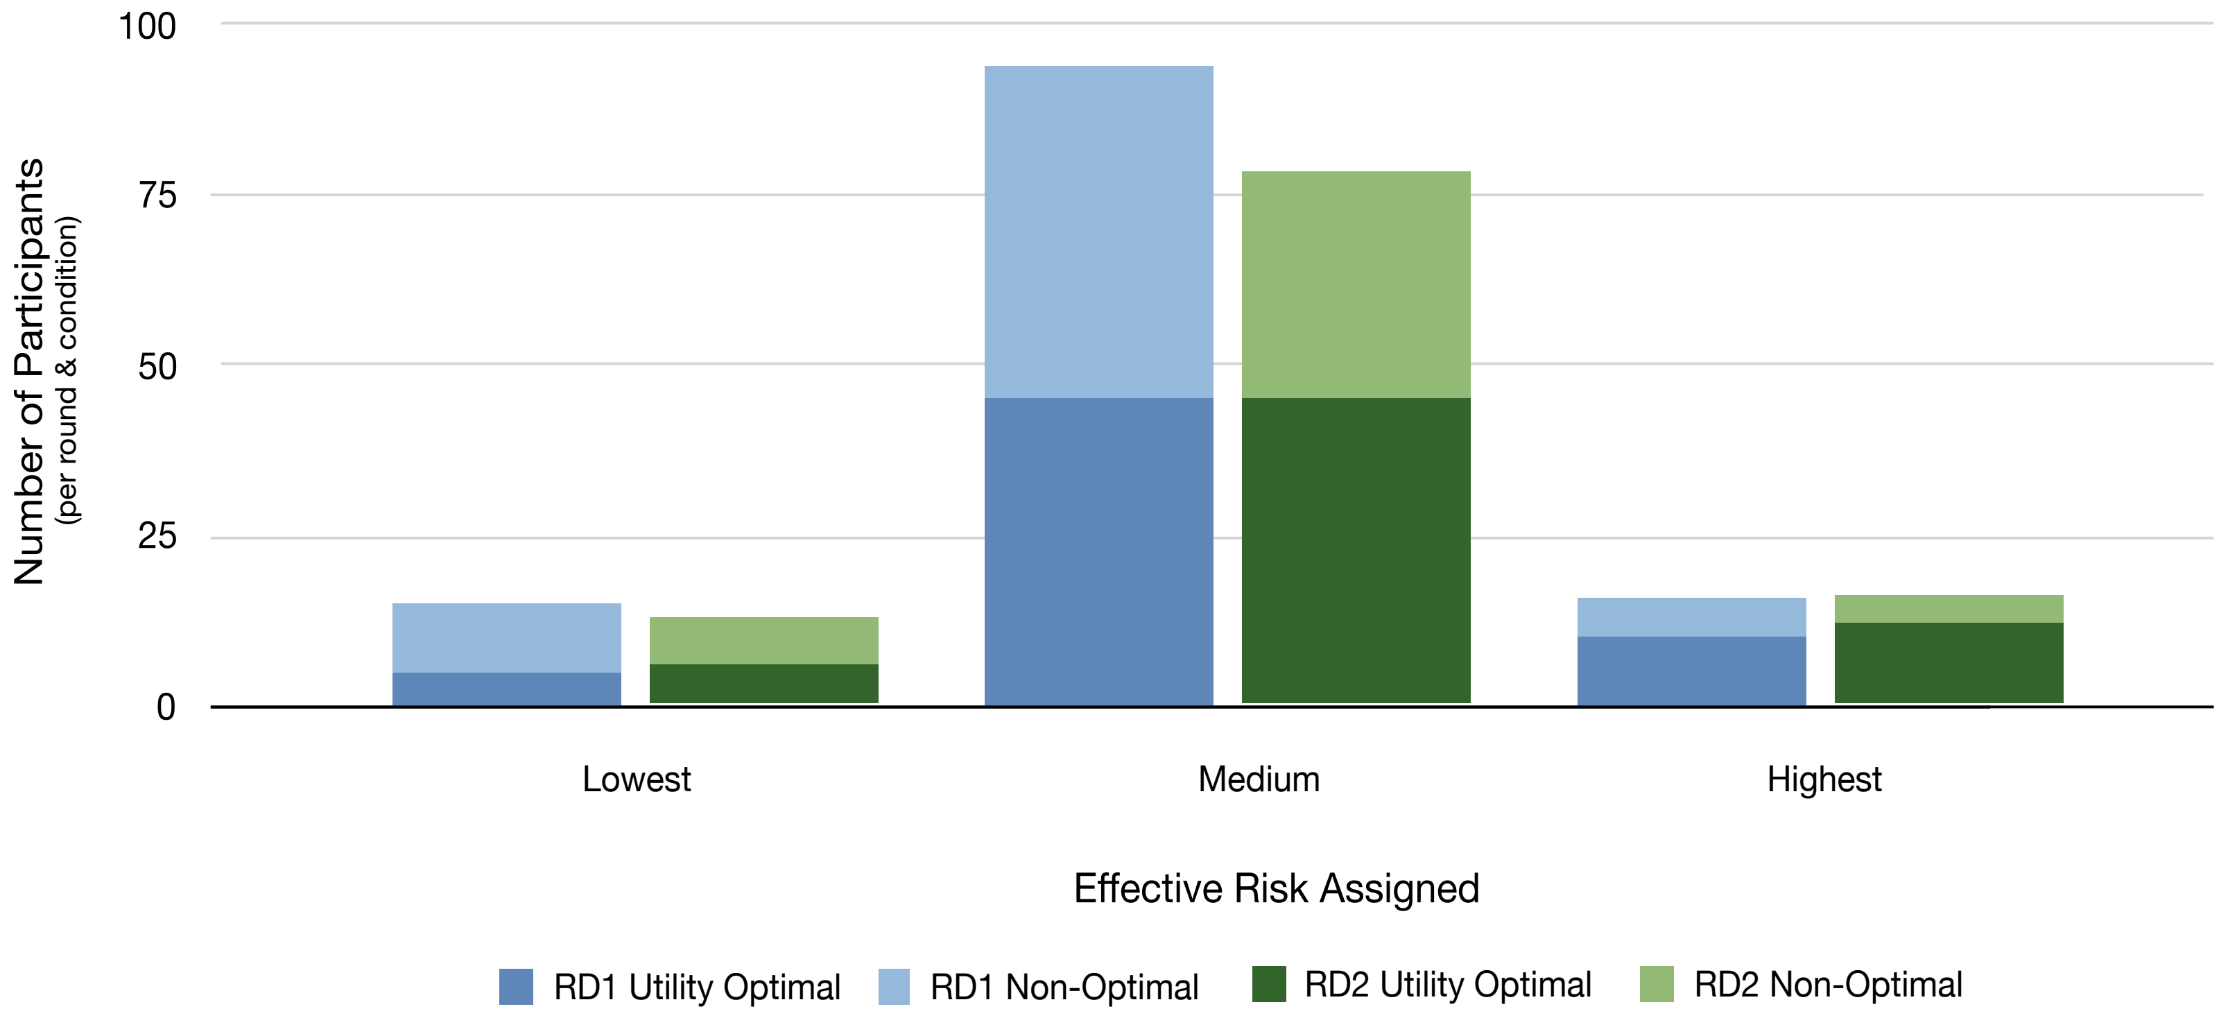
Number of Participants
(per round & condition)
Lowest
Medium
Highest
Effective Risk Assigned
RD1 Utility Optimal
RD1 Non-Optimal
RD2 Utility Optimal
RD2 Non-Optimal