

Picking vs. Guessing Secrets: A Game-Theoretic Analysis

MHR Khouzani*, Piotr Mardziel†, Carlos Cid‡, Mudhakar Srivatsa§

*Queen Mary, University of London

†University of Maryland, College Park

‡Royal Holloway, University of London

§IBM T.J. Watson Research Laboratory

Abstract—Choosing a hard-to-guess secret is a prerequisite in many security applications. Whether it is a password for user authentication or a secret key for a cryptographic primitive, picking it requires the user to trade-off usability costs with resistance against an adversary: a simple password is easier to remember but is also easier to guess; likewise, a shorter cryptographic key may require fewer computational and storage resources but it is also easier to attack. A fundamental question is how one can optimally resolve this trade-off. A big challenge is the fact that an adversary can also utilize the knowledge of such usability vs. security trade-offs to strengthen its attack.

In this paper, we propose a game-theoretic framework for analyzing the optimal trade-offs in the face of strategic adversaries. We consider two types of adversaries: those limited in their number of tries, and those that are ruled by the cost of making individual guesses. For each type, we derive the mutually-optimal decisions as Nash Equilibria, the strategically pessimistic decisions as maximin, and optimal commitments as Strong Stackelberg Equilibria of the game. We establish that when the adversaries are faced with a capped number of guesses, the user’s optimal trade-off is a uniform randomization over a subset of the secret domain. On the other hand, when the attacker strategy is ruled by the cost of making individual guesses, Nash Equilibria may completely fail to provide the user with any level of security, signifying the crucial role of credible commitment for such cases. We illustrate our results using numerical examples based on real-world samples and discuss some policy implications of our work.

Index Terms—Password Attacks; Attacker-Defender Games; Usability-Security Trade-off; Game Theory; Decision Theory; Maximin; Nash Equilibrium; Strong Stackelberg Equilibrium.

I. INTRODUCTION

Passwords remain the most common means of authenticating humans to computer systems. Yet, passwords are also among the most common points of failure of security systems [1–3]. According to an investigation report in 2011 [4], stolen login credentials accounted for nearly a third of corporate data breach incidents, out of which, more than a quarter were estimated to be carried out using a form of a guessing attack. Poorly chosen passwords undermine an otherwise secure authentication system. Users tend to choose easy to remember passwords [5, 6]. This is rationalizable as attempts in balancing usability costs with perceived security.

Managing the utility vs. security trade-off is also a relevant problem in the application of cryptographic techniques, which usually rely on maintaining a key unknown to any adversary.

Longer keys provide stronger security guarantees but at the same time inflict larger storage and computational costs on the system. Using cryptographic techniques, therefore, entails trading off utility for security, either through the choice of key size or the method of key generation. Either way, the decision must be made in the context of adversaries.

Guessing attacks are often categorized as *online* and *offline* based on their context of execution [7]. Online attacks involve interacting with the target system. In such an attack, adversaries are often limited in the number of (failed) guesses they can make (within a certain time period) before the system prevents any further interaction. In the case of password authentication, this is usually an account lock-out that requires intervention of the legitimate user using an alternate channel of authentication (email, phone, etc.).

In *offline* attacks, adversaries are assumed to have collected sufficient data to examine unlimited number of guesses, and are only constrained by their computational resources. In the case of password authentication, for example, this data can be the leaked hashes of user passwords, enabling the attackers to compute hashes of their guesses and compare them for a match, theoretically an unlimited number of times. Another example of an offline attack setting is when an adversary eavesdrops a cryptographic response to a predictable challenge in a challenge-response authentication protocol. Although unlimited in the number of guesses, adversaries in such offline scenarios still need to be wary of costs of trying guesses as computation of password hashes or cryptographic responses are not instantaneous or free (specially, noting that hash functions for hashing passwords are intentionally chosen to be slow on hardware to dissuade brute-force attacks). Hence, the response of such adversaries is governed by the computational/time cost per each guess. An adversary may obtain a pre-computed list of hashes to remove (or a *rainbow* table to mitigate) the computational burden during the execution of the attack. In such cases, the bottle-neck becomes the storage requirement for such a table, which implies a cap on the number of available guesses, similar to the online case.

We will collectively refer to passwords or cryptographic keys as *secrets*. We also use the terms *Capped-Guesses* and *Costly-Guesses* to respectively describe the following two settings: (1) adversaries are limited in their number of guesses, e.g. in online password attacks in the presence of a rate

limiting mechanism, or in offline attacks that use storage-limited pre-computed tables; and (2) adversaries incur a cost per each guess, e.g. in brute-force offline attacks. Regardless of the type of the guessing attack, the inherent behavioral or systematic preferences over the secret space can be exploited by adversaries and boost their guessing efficiency. Therefore, any secret picking policy that aims to achieve a desirable trade-off between usability and security must evaluate the possible reaction of a rational adversary given their capabilities. In particular, it is insufficient to analyze the decisions of either the users or the adversaries without taking into account the reaction of the other. Game theory provides tools to analyze such strategic interactions. The notion of *equilibrium*, in particular, describes how rational parties would eventually behave when faced against each other by characterizing their mutually-optimal strategies.

The basic question at the heart of this paper is the following: given a known uneven usability cost over the space of secrets, how can the defender optimally randomize in picking a secret? The main contribution of the paper is answering this fundamental question. Specially:

- We present novel decision and game-theoretic models for both Capped-Guesses and Costly-Guesses settings that are simple enough to allow analysis yet general enough to cover all the cases described above.
- We provide complete analysis of these games and discuss the security implications of the solutions. Specifically, we derive optimal secret selection policies with respect to different strategic metrics, namely, the strategically pessimistic solutions (Maximin), the mutual-best-response solutions (Nash Equilibria – NE), and the optimal commitment strategies (Strong Stackelberg Equilibria – SSE).
- For Capped-Guesses settings, we show that, interestingly, the optimal picking strategies still constitute uniform distributions despite the uneven preferences of the picker over the secret space. The trade-off is achieved by randomizing only over a (lower cost) subset of the secret space, while the probability distribution over the subset is uniform. The size of the subset is influenced by the picker’s trade-off parameters and (only) the cap on the available guesses. The optimal guessing strategies are restricted to the same subset though they are not uniform. Instead, the guesser probes the picker’s more favored secrets in that subset with higher probabilities. We also show that for this scenario, all of the different strategic metrics of Maximin, NE and SSE lead to the same solution for the picker.
- For the Costly-Guesses settings, we find a surprising result, reminiscent of the prisoner’s dilemma situation: aside from trivial cases, the NE strategies of the picker fail to yield any desirable security level, irrespective of the size of the secret space or the cost associated with the loss of the secret. We demonstrate how the picker can retrieve a desirable usability-security trade-off using *commitment* to optimal randomizations. We also notice that these optimal commitment (SSE) strategies for this case are almost never completely uniform, though they resemble uniform selec-

tion, with diminishing tails on costlier secrets.

- We provide numerical illustrations of our analyses using examples such as the leaked RockYou password dataset and cryptographic keys with increasing costs in their size.

The paper is structured as follows: Sec. II introduces the building blocks of our *non-zero-sum two-player* game between a picker and a guesser. In Sec. III, we present the model for Capped-Guesses scenarios and introduce different game theoretic notions of a solution, which we fully derive in Sec. IV. In Sec.V and VI, we present the model and analysis of the Costly-Guesses scenarios. In Sec. VII we comment on some of the implications of our results. A brief overview of related literature is discussed in Sec. VIII. A summary of our results and some suggestions for future directions of research concludes our paper in Sec. IX. Most of the technical proofs in the paper are relegated to the appendices in our technical report[8].

II. MODEL

In what follows, we progressively construct the model of our non-zero-sum two-person games between the *picker*, and the *guesser*. Critically, we assume that the parameters of the games are “common knowledge”, i.e., both players are aware of the presence and type of the game, the utilities and the information available to each other.

The picker (she) chooses a *secret* from the finite set of all secrets $\mathcal{P} = \{p_1, \dots, p_{|\mathcal{P}|}\}$. Let $d \in \mathcal{P}$ denote a pure (i.e., deterministic) *action* of the picker. \mathcal{P} is thus the picker’s *pure action set*. The picker has uneven preferences over this set of secrets. In the case of password selection, for instance, this preference could be related to the memorability and ease of use: simpler passwords are easier to remember and less cumbersome to type in. In the spirit of the von Neumann-Morgenstern utility theorem [9], we model these preferences by assigning different costs to different secrets.¹ Specifically, let the whole set of secrets be partitioned into disjoint non-empty subsets $\mathcal{E}_1, \dots, \mathcal{E}_N$, i.e., $\mathcal{E}_i \neq \emptyset$ for all i , $\mathcal{E}_i \cap \mathcal{E}_j = \emptyset$ for $i \neq j$ and $\cup_{i=1}^N \mathcal{E}_i = \mathcal{P}$, such that the picker incurs a usability cost of C_i if she picks any of the members of the set \mathcal{E}_i as her secret. Without loss of generality, assume $0 \leq C_1 < \dots < C_N$. Hence, in the absence of any other considerations, the picker prefers to choose her secret from set \mathcal{E}_i rather than \mathcal{E}_j when $i < j$, as she assigns a lower usability cost to secrets from the first set. These data are determined, for instance in the case of password choice, by statistical investigation of the past databases of cracked passwords, e.g., as published in [10, 11]. Alternatively, these sets can represent passwords that minimally satisfy an *increasingly more complex* password creation rule-sets. For instance, \mathcal{E}_1 can be the set of all dictionary words in lower case, \mathcal{E}_2 the set of all dictionary words but requiring a mix of capital letters, \mathcal{E}_3 having the additional rule of including a number as well, \mathcal{E}_4 requiring a

¹Note, however, that we assume the usability costs and security costs of the picker are additive through an appropriate scaling.




Fig. 1. (Top) RockYou dataset password frequency and the derived cost model. (Bottom) Fraction of passwords guessed as a function of number of guesses.

symbolic character too, etc.²

The guesser (he) makes guesses about the choice of the picker. Upon the discovery of the secret, i.e., a correct guess, the guesser wins a gain of $\gamma > 0$, and the picker incurs a loss of $\lambda > 0$. The guesses are either constrained in number or subject to cost. We will investigate these two cases separately in Sections III and V respectively. In what follows, we provide two numerical instances of the model. Note that these numerical examples are mainly for the purpose of illustrating the analysis.

A. Cost Example: Passwords

The RockYou password dataset [12] contains the passwords of around 32 million users of the RockYou gaming site. The data-breach that produced the list was particularly costly as the site did not bother hashing its users' passwords. The list is complete, containing both very common passwords (the password "123456" occurs 290729 times), as well as many unique ones (2459760 passwords appear only once). As a result, the list has been studied extensively [2, 10, 13–15]. Fig. 1 (top) summarizes the frequency (dark line) of the passwords in the whole dataset. The passwords in the figure are ordered in decreasing frequency of appearance.

The dashed line in Fig. 1 (bottom) demonstrates the strength of the passwords in the dataset using a simple metric quantifying the likelihood of a successful brute-force attack against a uniformly picked user in the dataset as a function of number of guesses, assuming the attacker knew the exact distribution of passwords in the dataset. As a frame of reference, we

²It is natural to assume that these partitions are common knowledge as *sets* as opposed to *lists*. In particular, no specific indexing of the members of a partition is common knowledge, and hence, the solutions must be *symmetric within the partitions*. Nevertheless, we provide our analysis agnostic of any assumption about existence of a common indexing inside partitions. Symmetric solutions can then be extracted.




Fig. 2. Synthetic cost model for selection of cryptographic keys, with cost proportional to linear and cubic power of key length.

also include a similar metric assuming the users picked their password uniformly from the 11884632 different passwords in the dataset (solid dark line) or if all 32 million users picked their passwords uniquely (solid light line).

As a candidate for the partitions, we group the passwords based on their frequency of appearance as an indirect measure of their cost. Namely, we group the passwords with the highest frequency in \mathcal{E}_1 , passwords with the second highest frequency in \mathcal{E}_2 , so on, which makes the last partition \mathcal{E}_{2040} as the set of all the passwords that appear only once. We use the inverse of frequency of a password as a rough estimate of its usability cost. After normalization, we set the usability costs to range from $C_1 = 1/290729$ to $C_{2040} = 1$. The cost associated with each partition can be seen as the dotted line in Fig. 1 (top).

In Section IV-A and Section VI-B we will compare the behavior of users in the dataset to that of their equilibrium behavior for each of our two attack settings.

B. Cost Example: Cryptographic Keys

The selection of secret keys for cryptographic protocols is usually out of the hands of humans, nevertheless, the design decision of picking the strength of the key (usually a function of its length) entails the same cost/risk trade-offs.

Fig. 2 summarizes the space of possible keys for two hypothetical cryptographic constructs. In both we assume a key can be anywhere between 0 and 64 bits. The examples differ only in the costs associated with each key. A cost linear in the length of key approximately models the trade-off in symmetric key systems such as AES which uses a number of rounds proportional to key length. The cubic relation is more appropriate approximation of public/private schemes such as naive implementations of RSA whose computation time scales cubically with key length [16]. Note that in our analysis we will assume that the length of the key is not known to a guesser, something that is usually not true of public/private schemes.

In Sections IV-B, VI-C, VI-E we examine the equilibrium strategies of the picker and guesser given each of these two cost models.

III. CAPPED-GUESSES

In the *Capped-Guesses* scenario, the guesser, *without observing the action of the picker*, chooses at most K elements from the set of possible secrets \mathcal{P} , as his guesses. We assume,

naturally, that K does not depend on the actual guesses chosen. The pure action of the guesser, which we denote by A , is hence a subset of \mathcal{P} of size K , since it is in the guesser’s best interest to use all K of his guesses. In the case of password selection, for example, each action represents an instance of a pre-computed table with which the guesser chooses to launch a dictionary attack. The action set of the guesser is therefore: $\mathcal{A} := \{A | A \subset \mathcal{P}, |A| = K\}$, the set of all possible pre-computed tables of size K . The number K represents the prowess of the guesser determined by the physical limitations in place: for instance, in the case of pre-computed table attacks on passwords, K is determined by how much memory each hash entry occupies and how much total memory the attacker has available for the table. Alternatively, in an online password attack, it can be the number of tries he is allowed to make before getting locked out. We assume $K < |\mathcal{P}| = \sum_{i=1}^N |\mathcal{E}_i|$, since otherwise, the guesser trivially can find the secret with certainty. A (pure) *strategy profile* here is simply a pair of picker and guesser actions, $(d, A) \in \mathcal{P} \times \mathcal{A}$.

The problem is the following: determine best strategies for the picker to choose her secret and the guesser to construct his guessing dictionary, when both parties are *rational* decision makers. The problem can be modeled as a *simultaneous move* game. Note that in game-theory, the term “simultaneous move” does not necessarily imply synchronicity, rather, the lack of observation of the move of other players (or any signal about it) before making a move. Otherwise, there is a sequentiality in the occurrences of the actions taken in our problem: the picker picks first. In our Capped-Guesses game, the “actions” and “strategies” simply coincide. To complete the model, we next provide the utilities of the players given a strategy profile (d, A) . Let u_D and u_A represent the utilities of the picker and the guesser respectively. Compactly put, we have:

$$u_D(d, A) = -c(d) - \lambda \mathbf{1}_A(d), \quad u_A(d, A) = \gamma \mathbf{1}_A(d) \quad (1)$$

where $\mathbf{1}$ represents the indicator function,³ and $c(d) := \sum_{i=1}^N C_i \mathbf{1}_{\mathcal{E}_i}(d)$ is the usability cost of secret d . Note that this summation only contains one non-zero element: if the picked secret is from partition \mathcal{E}_i , she incurs the usability cost of C_i . A list of the main notations is provided in Table I.

Solution Concept 1 – Nash Equilibria (NE): A *solution* of a game is a prediction of how rational players facing it may take decisions. A commonly used notion of a solution is Nash Equilibrium (NE in short), which informally put, is a strategy profile that consists of simultaneously optimal responses to each other, keeping the others’ strategies fixed, i.e., strategy profiles that are resistant against unilateral deviations of players. Formally, in our two player game, this means the following: the strategy pair $(d^*, A^*) \in \mathcal{P} \times \mathcal{A}$ is a (pure) NE if and only if: $d^* \in \arg \max_{d \in \mathcal{P}} u_D(d, A^*)$ and $A^* \in \arg \max_{A \in \mathcal{A}} u_A(d^*, A)$, in (1), i.e., $u_D(d^*, A^*) \geq u_D(d, A^*)$ for all $d \in \mathcal{P}$ and $u_A(d^*, A^*) \geq u_A(d^*, A)$ for all $A \in \mathcal{A}$.

³The indicator (characteristic) function of a subset Y of a set X is a function $\mathbf{1}_Y : X \rightarrow \{0, 1\}$ defined as the following: $\mathbf{1}_Y(x) = 1$ if $x \in Y$, and $\mathbf{1}_Y(x) = 0$ if $x \notin Y$.

TABLE I
LIST OF MAIN NOTATIONS

| Notation | Definition |
|------------------|---|
| \mathcal{E}_i | Set of secrets with the same cost of picking C_i |
| \mathcal{P}, N | Set of all secrets, Number of its partitions |
| C_i | Cost of picking a secret from set \mathcal{E}_i , incurred by the picker |
| K | Number of attempts available to the guesser in Capped-Guesses (size of his “table”) |
| σ | Cost of each attempt of the guesser in Costly-Guesses |
| γ | Gain earned by the guesser if any of his guesses is correct |
| λ | Loss incurred by the picker if the secret is found by the guesser |
| $c(d)$ | Picker’s usability cost for picking secret d . Short for $\sum_{i=1}^N C_i \mathbf{1}_{\mathcal{E}_i}(d)$. |

It is not difficult to see that “pure” Nash Equilibrium is not a suitable solution concept for our game. In fact, except in trivial cases, no pure NE exists. This is because a pure strategy of the picker means selection of a specific secret. The best response of the guesser is then simply to include that secret in his guess dictionary. But then, the picker would have been better off to deviate and choose a different secret and not incur the potentially huge cost of having her secret revealed with certainty. In other words, deviation from any pure action (except in trivial cases) is beneficial for the picker. A similar argument can be made for the guesser.

The above discussion motivates the search for a solution among *mixed strategies*, which involve randomization thereby injecting ambiguity about the choice of each player. Specifically, a mixed strategy of a player is a probability distribution over her set of pure strategies. For any finite nonempty set \mathcal{S} , let $\Delta(\mathcal{S})$ represent the set of all probability distributions over it. That is:

$$\Delta(\mathcal{S}) := \{\sigma \in \mathbb{R}^{+|\mathcal{S}|} \mid \sum_{s \in \mathcal{S}} \sigma(s) = 1\}$$

For a given probability distribution $\sigma \in \Delta(\mathcal{S})$, let the *support* of σ , or $\text{supp}(\sigma)$, denote the subset of the domain of σ that receives a strictly positive probability, that is:

$$\text{supp}(\sigma) := \{s \in \mathcal{S} \mid \sigma(s) > 0\}.$$

Moreover, for a given probability distribution $\sigma \in \Delta(\mathcal{S})$ and a given subset $\mathcal{S}' \subseteq \mathcal{S}$, let $\sigma(\mathcal{S}')$ represent the probability measure of \mathcal{S}' with respect to σ , that is, let $\sigma(\mathcal{S}') := \sum_{s \in \mathcal{S}'} \sigma(s)$.

Let δ and α represent a mixed strategy of the picker and guesser respectively. We hence have: $\delta \in \Delta(\mathcal{P})$ and $\alpha \in \Delta(\mathcal{A})$. Following a common abuse of notation in game theory, let $u_D(\delta, \alpha)$ and $u_A(\delta, \alpha)$ be the expected utility of the two players given a mixed strategy profile $(\delta, \alpha) \in \Delta(\mathcal{P}) \times \Delta(\mathcal{A})$ where the expectation is taken with respect to the independent randomizations in the mixed strategies. That is: $u_D(\delta, \alpha) := \sum_{d \in \mathcal{P}, A \in \mathcal{A}} u_D(d, A) \delta(d) \alpha(A)$, and likewise

for $u_A(\delta, \alpha)$. Replacing from (1), we have:

$$u_D(\delta, \alpha) = - \sum_{d \in \mathcal{P}} c(d) \delta(d) - \lambda \sum_{d \in \mathcal{P}, A \in \mathcal{A}} \mathbf{1}_A(d) \delta(d) \alpha(A) \quad (2a)$$

$$u_A(\delta, \alpha) = \gamma \sum_{d \in \mathcal{P}, A \in \mathcal{A}} \mathbf{1}_A(d) \delta(d) \alpha(A) \quad (2b)$$

Note that we are assuming randomization per se is costless. A mixed strategy of the guesser, $\alpha \in \Delta(\mathcal{A})$, specifies the probability that each feasible dictionary (table) is selected. For our model, it is often simpler to instead specify the marginal probabilities that each secret is tried by the guesser. Specifically, let us define ρ such that $\rho(d)$ denotes the probability that secret d is in the (K -sized) table of the guesser. ρ and α are related through: $\rho(d) = \sum_{A \in \mathcal{A}} \mathbf{1}_A(d) \alpha(A)$. Moreover, using the notion of probability measure and the fact that all members of the same partition by definition have the same choosing cost for the picker, we have: $\sum_{d \in \mathcal{P}} c(d) \delta(d) = \sum_{i=1}^N C_i \delta(\mathcal{E}_i)$. Hence, the expressions in (2) can be simplified as:

$$u_D(\delta, \rho) = - \sum_{i=1}^N C_i \delta(\mathcal{E}_i) - \lambda \sum_{d \in \mathcal{P}} \delta(d) \rho(d), \quad (3)$$

$$u_A(\delta, \rho) = \gamma \sum_{d \in \mathcal{P}} \delta(d) \rho(d)$$

A mixed NE is defined in the same way as a pure NE, except that the optimization variables and the optimization spaces are replaced accordingly. The set of pure NE are contained in the set of mixed NE, since pure strategies can be obtained from degenerate distributions over the strategies. That is, a mixed strategy profile (δ^*, α^*) is a mixed NE iff:

$$u_D(\delta^*, \alpha^*) \geq_{\forall \delta \in \Delta(\mathcal{P})} u_D(\delta, \alpha^*), \quad u_A(\delta^*, \alpha^*) \geq_{\forall \alpha \in \Delta(\mathcal{A})} u_A(\delta^*, \alpha).$$

Solution Concepts 2 & 3 – maximin and minimax: A (mixed) strategy of the picker $\delta^{\text{maximin}} \in \Delta(\mathcal{P})$ is a *maximin* strategy of hers if and only if:

$$\delta^{\text{maximin}} \in \arg \max_{\delta \in \Delta(\mathcal{P})} \left[\min_{\alpha \in \Delta(\mathcal{A})} u_D(\delta, \alpha) \right]$$

Let $\underline{u}_D(\delta) := \min_{\alpha \in \Delta(\mathcal{A})} u_D(\delta, \alpha)$, which is the worst utility of the picker among all reactions of the guesser if she chooses the mixed strategy of δ . Then δ^{maximin} maximizes $\underline{u}_D(\delta)$, achieving the *maximin utility*, which we will denote by $\underline{u}_D^{\text{max}}$. This is the mixed strategy that guarantees (secures) the picker at least her maximin utility irrespective of the strategy of the guesser. For this reason, maximin strategies are sometimes also referred to as “security” strategies. maximin strategies are recipe for action when a player is strategically pessimistic, in that she believes the opponent(s) behave in such a way to hurt her utility the most, as opposed to selfishly maximize their own utilities. Hence, the focus is solely on the utility of that player, and rationality of other players is not taken into account.

This is conceptually different from a *minimax* strategy of a player. Formally, δ^{minimax} is a picker’s minimax strategy if and only if: $\delta^{\text{minimax}} \in \arg \min_{\delta \in \Delta(\mathcal{P})} [\max_{\alpha \in \Delta(\mathcal{A})} u_A(\delta, \alpha)]$. Let $\overline{u}_A(\delta) := \max_{\alpha \in \Delta(\mathcal{A})} u_A(\delta, \alpha)$, which is the best utility of

the guesser among all of his reactions if the picker chooses the mixed strategy of δ . Then δ^{minimax} minimizes $\overline{u}_A(\delta)$, guaranteeing that the utility of the guesser is bounded by his *minimax utility*, denoted by $\overline{u}_{A\text{min}}$. That is the strategy that the picker can adopt to hurt the utility of the opponent (the guesser) the most, ignoring her own utility. In zero-sum games, the utility of each player is negative (i.e., additive inverse) of the of other. Hence, hurting the expected pay-off of the opponent the most is exactly equivalent to helping your own expected pay-off the most. This means that minimax and maximin strategies of each of the players coincide. But this in general does *not* extend to non-zero-sum games. This is exactly the situation in our game. It is easy to see that the minimax strategy of the picker is simply to uniformly randomize over the *entire* set of secrets, effectively maximizing the ambiguity, minimizing any useful information that the guesser can exploit. However, this completely ignores the cost of choosing costly secrets. As we will show, the maximin strategy of the picker is in general different from uniform randomization over the entire set of secrets.

Likewise, we can speak of the maximin and minimax strategies of the guesser: $\alpha^{\text{maximin}} \in \Delta(\mathcal{A})$ is a maximin strategy of the guesser if and only if: $\alpha^{\text{maximin}} \in \arg \max_{\alpha \in \Delta(\mathcal{A})} \underline{u}_A(\alpha)$ where $\underline{u}_A(\alpha) := \min_{\delta \in \Delta(\mathcal{P})} u_A(\delta, \alpha)$. Here also the distinction between the maximin and minimax strategies can be observed. Specifically, if the guesser is on the (pessimistic) belief that the picker is trying to hurt his utility the most (or equivalently plan according to the “worst case scenario” of the strategy of the picker irrespective of her rationality), he should select his K guesses uniformly randomly over the entire set of secrets. This approach ignores the pay-off structure of the picker and hence does not take advantage of the presence of the preferences of the picker over the secrets. We will see how the guesser can exploit this knowledge in Sec. IV.

Solution Concept 4 – Strong Stackelberg Equilibria (SSE): Consider the situation in which the picker has the power of credible commitment to a mixed strategy. Note that this is in general different from commitment to a pure strategy and requires a different “apparatus”. The relevant solution concept for these cases is the Strong Stackelberg Equilibria, which intuitively put, are the best mixed strategies that the leader (picker in our case) can commit to, knowing that the follower (guesser, here) will observe this commitment and will respond selfishly optimally to it. In order for the solution concept to exist, it also needs the extra assumption that whenever the follower is indifferent between a set of best responses, he will break ties in favor of the leader. This is a benign assumption, because the leader can turn any of the indifferent best responses of the follower to a strict preference through an infinitesimal modification of her mixed strategy. Note that a (pure) strategy of the follower is now a function of the commitment distribution of the leader. That is, if the follower is the guesser, a pure strategy of the follower is a mapping from $\Delta(\mathcal{P})$ to \mathcal{A} . Formally, $(\delta^*, \alpha^{\text{BR}})$ in which $\delta^* \in \Delta(\mathcal{P})$ and $\alpha^{\text{BR}} : \Delta(\mathcal{P}) \rightarrow \Delta(\mathcal{A})$, constitutes a SSE

Game 1: Capped-Guesses

Players: PICKER, GUESSER

Strategy Sets: PICKER'S: $\{d \in \mathcal{P}\}$

GUESSER'S: $\{A \subset \mathcal{P}, |A| = K\}$

Utilities: PICKER: $u_D(d, A) = -c(d) - \lambda \mathbf{1}_A(d)$,

GUESSER: $u_A(d, A) = \gamma \mathbf{1}_A(d)$

if and only if:⁴

- 1) $\delta^* \in \arg \max_{\delta \in \Delta(\mathcal{P})} u_D(\delta, \alpha^{\text{BR}}(\delta))$
- 2) $\alpha^{\text{BR}}(\delta) \in \arg \max_{\alpha \in \Delta(\mathcal{A})} u_A(\delta, \alpha)$
- 3) $\alpha^{\text{BR}}(\delta) \in \arg \max_{\alpha' \in \arg \max_{\alpha \in \Delta(\mathcal{A})} u_A(\delta, \alpha)} u_D(\delta, \alpha')$

IV. ANALYSIS OF THE CAPPED-GUESSES SCENARIO

As our main result for the Capped-Guesses scenario, we provide a sufficient condition for a strategy pair to be a mixed NE (Prop. 1). We show that the NE and maximin strategies of the picker coincide (Lemma 1). This useful property leads us to other implication: all NE are interchangeable (Corollary 1) and they all yield the same utility for the picker (Corollary 2). Another implication of the lemma is that for this scenario, the set of optimal mixed strategies of the picker to commit to, i.e., her SSE strategies, are also the same as her NE strategies, and moreover, they attain her the same utility as any NE does (Corollary 3). Finally, we provide a mild constraint under which the sufficient conditions provided in Prop. 1 for a mixed strategy of the picker to be a NE are also necessary conditions, implying uniqueness of the description of the NE for almost all instances of the game (Corollary 6). These results fully characterize the solution of the Capped-Guesses game. The proofs of the results in this section can be found in the Appendices of our accompanying technical report[8].

First, note that following Nash's Theorem, our finite game has at least one mixed NE. The existence of maximin, minimax and SSE solutions also follow standard results in game theory [17]. In order to explicitly describe the NE, we need to define a few parameters. Let: $L := \min_{1 \leq l \leq N} l$ s.t. $\sum_{i=1}^l |\mathcal{E}_i| > K$. Note that in part this means: $|\cup_{i=1}^m \mathcal{E}_i| \leq K$ for any $m < L$ (recall that K is the dictionary size of the guesser – the available number of guesses to the adversary). Now suppose the picker chooses her secret according to a randomization only from the first m (cheapest) partitions where $m < L$. Then the guesser can correctly guess the secret with certainty, because he can simply include the entire $\cup_{i=1}^m \mathcal{E}_i$ in his guessing dictionary. Hence, for the picker, the (strictly) best among such options that lead to certain loss of the secret is simply picking from the cheapest partition which yield her a utility of $-C_1 - \lambda$.⁵ The picker can reduce the chance of a correct guess by randomizing over partitions beyond $\cup_{i=1}^m \mathcal{E}_i$,

⁴The superscript BR is chosen to stand for “best response”.

⁵In the language of game theory, any mixed strategy of the picker that only randomizes over $\cup_{i=1}^m \mathcal{E}_i$ where $m < L$ is strictly dominated by strategies that only randomize over \mathcal{E}_1 .

but then the picker has to balance usability costs with the gain in increasing the entropy. Define:

$$\mathcal{J} := \left\{ L < j \leq N \mid \lambda K + \sum_{i=1}^{j-1} C_i |\mathcal{E}_i| \geq C_j \sum_{i=1}^{j-1} |\mathcal{E}_i| \right\}. \quad (4)$$

That is, \mathcal{J} characterizes the partitions for which the inequality of $\lambda K / (\sum_{i=1}^{j-1} |\mathcal{E}_i|) + (\sum_{i=1}^{j-1} C_i |\mathcal{E}_i|) / (\sum_{i=1}^{j-1} |\mathcal{E}_i|) \geq C_j$ holds. Since only $j > L$ are considered, we have $K < (\sum_{i=1}^{j-1} |\mathcal{E}_i|)$. In particular, suppose the picker *uniformly* randomizes over $\text{unif}(\cup_{i=1}^{j-1} \mathcal{E}_i)$. Then, irrespective of the strategy of the guesser as long as its support is $\text{unif}(\cup_{i=1}^{j-1} \mathcal{E}_i)$, his chance of finding the secret is exactly $K / (\sum_{i=1}^{j-1} |\mathcal{E}_i|)$, and hence the security cost of the picker is $\lambda K / (\sum_{i=1}^{j-1} |\mathcal{E}_i|)$. Moreover, the usability cost of the picker for uniformly randomizing over $\text{unif}(\cup_{i=1}^{j-1} \mathcal{E}_i)$ is $(\sum_{i=1}^{j-1} C_i |\mathcal{E}_i|) / (\sum_{i=1}^{j-1} |\mathcal{E}_i|)$. Therefore, the condition in the definition of \mathcal{J} translates to the following: $j \in \mathcal{J}$ if the usability cost of choosing from \mathcal{E}_j is less than the overall cost (security and usability cost) of uniformly randomizing over the (combined) first $j - 1$ (cheapest) partitions.⁶

If $\mathcal{J} \neq \emptyset$, define $J := \max \mathcal{J}$. We label the cases where either $\mathcal{J} = \emptyset$ or $C_1 + \lambda \leq (\sum_{i=1}^J C_i |\mathcal{E}_i| + \lambda K) / (\sum_{i=1}^J |\mathcal{E}_i|)$ as “total defeat”, since in such cases the picker chooses her secret from the cheapest partition, \mathcal{E}_1 , knowing that her choice will be guessed correctly, because it is not worthwhile (or not possible) for her to try to prevent it. We will refer to all other situations, i.e., when we have $\mathcal{J} \neq \emptyset$ and $C_1 + \lambda > (\sum_{i=1}^J C_i |\mathcal{E}_i| + \lambda K) / (\sum_{i=1}^J |\mathcal{E}_i|)$ as “ordinary” cases, since, as we show, it is worthwhile for the picker to try to avoid certain revelation of her secret.

Recall that $\rho^*(p) = \sum_{A \in \mathcal{A}} \alpha^*(A) \mathbf{1}_A(p)$ is just the probability that secret p will be among the K selections of the guesser, given his mixed strategy of α^* . We now mathematically present the NE strategies and subsequently describe them in words:

Proposition 1: For the “ordinary” cases in a Capped-Guesses game, consider a strategy pair (δ^*, α^*) where:

$$\delta^* = \text{unif}(\cup_{i=1}^J \mathcal{E}_i),$$

and:

$$\rho^*(p) = \frac{K}{\sum_{j=1}^J |\mathcal{E}_j|} + B_i, \quad \forall p \in \mathcal{E}_i \text{ where } i \leq J \quad (5a)$$

$$\rho^*(p) = 0, \quad \forall p \in \mathcal{E}_i \text{ where } i > J \quad (5b)$$

where $B_i := \left[\sum_{j=1}^J C_j |\mathcal{E}_j| - C_i \sum_{j=1}^J |\mathcal{E}_j| \right] / \left[\lambda \sum_{j=1}^J |\mathcal{E}_j| \right]$. Then, the strategy pair (δ^*, α^*) is a (mixed) NE. For the “total

⁶With simple algebra, the condition can be shown to be equivalent to the following: $[\lambda K / (\sum_{i=1}^{j-1} |\mathcal{E}_i|) + (\sum_{i=1}^{j-1} C_i |\mathcal{E}_i|) / (\sum_{i=1}^{j-1} |\mathcal{E}_i|)] \geq [\lambda K / (\sum_{i=1}^j |\mathcal{E}_i|) + (\sum_{i=1}^j C_i |\mathcal{E}_i|) / (\sum_{i=1}^j |\mathcal{E}_i|)]$. In words, $j \in \mathcal{J}$ if the overall cost of uniformly randomizing over the combined first $j - 1$ partitions is more than that of uniformly randomizing over the combined first j partitions for the picker. This in turn implies that, for the picker, uniform randomization over the first $j - 1$ partitions is (weakly) *dominated* by uniformly randomizing over the first j partitions.

defeat” cases, consider a strategy pair (δ^*, α^*) that satisfies the following:

$$\text{Picker: } \delta^*(\mathcal{E}_1) = 1 \quad (6)$$

$$\text{Guesser: } \begin{cases} \rho^*(p) > 1 - \frac{C_i - C_1}{\lambda} & \forall p \in \mathcal{E}_i, i \leq J, \\ \rho^*(p) = 0 & \forall p \in \mathcal{E}_i, i > J \end{cases} \quad (7)$$

Then (δ^*, α^*) constitutes a NE.

In words, for the “ordinary” cases, the proposed NE is the following: the picker chooses her secret *only* from the first J partitions, i.e., the J most favored partitions, and does so *uniformly* randomly. Note in particular that the preference profile of the picker only affects her NE strategy through the number of partitions that constitute the domain of secrets to choose from, but the randomization over this domain is always uniform, despite the uneven preferences over them.

On the other hand, the guesser, knowing the picker does not choose her secret with any positive probability from partitions beyond \mathcal{E}_J , does not include any guesses from them either (5b). The guesser selects uniformly randomly *within* partitions $1, \dots, J$ but *not across* them. That is, even though the secrets from the same partition are equally likely to be part of the guessing dictionary of the guesser, the secrets from partition $i \leq J$ are chosen with a *bias* equal to B_i away from uniform guessing. This is despite the fact that the picker chooses her secret uniformly randomly from the first J partitions. Indeed, as we discuss in the proof, the guesser explores the relatively favored partitions of the picker among the first J partitions with a positive bias compared to her relatively less favored partitions. Specifically, the bias is exactly such that the picker is indifferent about choosing the secret from any of the first J partitions.

For the cases of “total defeat”, the picker simply chooses her secret from partition \mathcal{E}_1 , the least costly partition, and the guesser includes all of that partition into his dictionary, along with other partitions such that the picker is forced into picking her secret only from the cheapest partition. Thus, the secret will be discovered by the guesser with probability one. Note that, interestingly, the NE was not at all affected by γ , the gain parameter of the guesser.

Our next series of results describe the properties of the NE in regards to other strategic metrics. Note that establishing these results do not rely directly on the explicit expression of the NE in Prop. 1.

In general, playing NE strategies by a player conjures the assumption that the other player(s) are indeed *rational*, in that, they are interested in maximizing their own utility as opposed to antagonistically trying to minimize the utility of that player. But what if this rationality assumption cannot be made in our case regarding the guesser? Our next observation dispels that concern by establishing that for the Capped-Guess scenarios, NE strategies of the picker are her maximin strategies and vice versa.

Lemma 1: Let $\Omega_{\text{NE}}^{\text{D}}$ be the set of NE strategies of the picker and $\Omega_{\text{maximin}}^{\text{D}}$ be the set of her maximin strategies in a game of Capped-Guesses. We have: $\Omega_{\text{NE}}^{\text{D}} = \Omega_{\text{maximin}}^{\text{D}}$.

The lemma establishes that the picker can randomize according to her NE and (in expectation) be guaranteed at least the expected utility prescribed by the NE, irrespective of the mixed strategy of the guesser, be it a NE or not. From a different viewpoint, the picker can act according to her pessimistic maximin strategy, but be assured that she does not lose anything in expectation by not playing a NE. Note that this property only holds for the NE strategy of the picker and not of the guesser (Recall that the maximin strategy of the picker is choosing his K guesses uniformly randomly from the entire secret space \mathcal{P}).

Here, we just mention the gist of the proof.⁷The argument starts by noting from (2) that for any $\delta \in \Delta(\mathcal{P})$, $\alpha^* \in \arg \max_{\alpha \in \Delta(\mathcal{A})} u_A(\delta, \alpha)$ if and only if: $\alpha^* \in \arg \min_{\alpha \in \Delta(\mathcal{A})} u_D(\delta, \alpha)$. To see this, note that the pay-off of the picker is composed of two parts, the first part is the expected cost of choosing the secret, and the second part is the expected cost of losing it. For any given mixed strategy of the picker, the guesser can only affect the second part of the utility of the picker. Specifically, $u_D(\delta, \alpha) = -(\lambda/\gamma)u_A(\delta, \alpha) + \phi(\delta)$, where $-(\lambda/\gamma) < 0$ and $\phi(\delta)$ is an expression that does not depend on α . That is, the (rational) best response of the guesser to any “given” strategy of the picker, also yields the worst utility for the picker. Hence, assuming a rational best response and strategically worst case scenario become equivalent for the picker.

Next two results (corollaries of Lemma 1) establish the interchangeability of the NE and remove the concern of “Equilibrium Selection” in games of Capped-Guesses.

Corollary 1: Interchangeability of NE (I): If (δ^*_1, α^*_1) and (δ^*_2, α^*_2) are both NE in a game of Capped-Guesses, then so are (δ^*_1, α^*_2) and (δ^*_2, α^*_1) .

This corollary shows that if at all there are more than one distinct NE present, then no matter which NE strategy each player chooses to play, the outcome is still a NE. The next corollary further shows that, even if there were multiple NE, there is no question of preference between them for the picker, since her utility is the same in all of them:

Corollary 2: Interchangeability of NE (II): All NE in a Capped-Guesses game yield the same utility for the picker. Specifically, if (δ^*, α^*) is a NE of the Capped-Guesses game, then: $u_D(\delta^*, \alpha^*) = \underline{u_D}^{\text{max}}$.

These two results imply that, as far as the picker is concerned, it suffices to to find “a” NE, as we did in Prop. (1), which is in general easier than finding the set of all NE. Although in our game, we will show that, almost in all cases, the NE is in fact unique (Corollary 6). The next corollary states that a NE strategy of the picker is also an optimum strategy of her to commit to, and vice versa.

Corollary 3: In a Capped-Guesses game, let $\Omega_{\text{SSE}}^{\text{D}}$ be the set of picker’s SSE strategies. Then: $\Omega_{\text{SSE}}^{\text{D}} = \Omega_{\text{NE}}^{\text{D}}$.

As in Lemma 1, the corollary follows by showing that given the committed strategy of the picker, the guesser will try to maximize his own utility, which in our Capped-Guesses

⁷As a reminder, the detail of this, as well as all the other proofs in this manuscript, are available in our technical report [8].

game, is exactly what he would do if he wanted to minimize the utility of the picker. Hence the best mixed strategy to commit to by the picker is exactly the strategy that maximizes her minimum utility, i.e., her maximin strategy, which we previously showed to match the NE strategies. Intuitively, this is because in the Capped-Guesses model, the guesser will enter the game irrespective of the randomization strategy of the picker, and use all of his K attempts. Moreover, he chooses his K guesses so as to maximize the chances of finding the secret, which is exactly antagonistic to the utility of the picker given the randomized strategy of the picker (refer to the discussion after Lemma 1).

Note that the ability to commit to a mixed strategy is guaranteed not to hurt the “committer” (leader), since the leader can always commit to her Nash strategies and yield at least her Nash utilities [18]. Or commit to a maximin strategy and guarantee her maximin utility. But in general, she may be able to do better and improve upon her Nash Equilibria. Even in the presence of Corollary 3, due to the property that in SSE, the follower breaks ties among his best responses in favor of the leader, identical SSE and NE strategies of the picker may lead to distinct utilities for her. However, the following lemma establishes that for the game of Capped-Guesses, this is not the case: the power to commit does not “buy” the picker any extra benefit. Specifically, the utility of the picker when best-committing is no better than her maximin utility.

Corollary 4: Let $(\delta^*, \alpha^{\text{BR}})$ be a SSE of a Capped-Guesses game. Then we have: $u_D(\delta^*, \alpha^{\text{BR}}(\delta^*)) = u_D^{\text{max}}$.

This result can also be expressed in the measure of the “value of mixed commitment” as discussed in [19]: the value of mixed commitment for the picker in Capped-Guesses games is one, i.e., commitment achieves nothing above what is achievable in NE, and hence there is no advantage in commitment. As we will see in Section VI, this is *drastically* different from the situation in Costly-Guesses scenarios.

Corollary 5: The NE strategies of the picker as described in Prop. 1 are also maximin and SSE. Moreover, her utility in all NE and SSE is her maximin utility, given as: $u_D^{\text{max}} = -\left[\sum_{j=1}^J C_j |\mathcal{E}_j| + \lambda K\right] / \left[\sum_{j=1}^J |\mathcal{E}_j|\right]$ in “ordinary” cases, and $u_D^{\text{max}} = -C_1 - \lambda$ in “total defeat” cases.

The next corollary is rather less important in characterization of this game in the light of Corollary 2. Nevertheless, it also shows that not only the utilities, but in fact even the equilibrium strategies themselves are almost always unique. This removes the question whether there may be other simpler to play NE of the game than presented in Prop. 1 (even though the NE for the picker is quite simple as is). The answer is no, almost never. Referring to the definition of J in (4), it allows to have $\left(\sum_{j=1}^{J-1} C_j |\mathcal{E}_j| + \lambda K\right) = C_J \left(\sum_{j=1}^{J-1} |\mathcal{E}_j|\right)$. We will refer to such a case as a *degenerate case*, which is completely identifiable from the parameters of the problem. For all other (“non-degenerate”) cases, the condition $\left(\sum_{j=1}^{J-1} C_j |\mathcal{E}_j| + \lambda K\right) \geq C_J \left(\sum_{j=1}^{J-1} |\mathcal{E}_j|\right)$ is *strictly* satisfied.

Corollary 6: Aside from degenerate cases identified above,




Fig. 3. (Top) Picker loss (solid) and corresponding key space size (dotted) as a function of the number of available guesses to the adversary for different values of loss (λ). (Bottom) Chance of a correct password guess as a function of the number of guesses in RockYou dataset and for a range of loss (λ) values.

the sufficient conditions for a NE strategy profile provided in Prop. 1 are also necessary.

Note that the corollary in part implies that for “non-degenerate” “ordinary” cases, the NE is unique.

A. Equilibrium Example: Passwords

Fig. 3 summarizes the equilibrium behaviors of the picker engaged in the Capped-Guesses game. The top part of the figure shows picker loss (negative of the utility) as solid lines and the size of the support set over which the picker chooses his passwords, as the dotted lines. The bottom part of the figure shows the probability that the password would be found by the guesser. All of these are shown as functions of the number of guesses available to the guesser and for two different values of λ . Recall that the cost of picking passwords was normalized. In this manner λ serves as the cost of security losses for having the password guessed relative to the usability cost.

As the adversary is granted more guesses, the picker has to include a larger subset of passwords to (uniformly) randomize over. When this support set is exhausted or the additional cost exceeds the benefits, the picker gives up and picks only the cheapest password (“total defeat”). For high values of λ , the picker never gives up, specifically, for large enough number of the available guesses, the picker uniformly randomizes over the entire set of passwords.

B. Equilibrium Example: Cryptographic Keys

Fig. 4 demonstrates the result of equilibrium key picking and key guessing. The top part of the figure shows the loss incurred by the picker in the two different cost models (linear or cubic in key size) as a function of the number of available guesses. The middle part of the figure focuses on the size of the key space that the picker is forced to choose from. The bottom




Fig. 4. Equilibrium in the Capped-Guesses model for key selection. (Top) picker loss as a function of number of available guesses. (Middle) the corresponding size of the support set. (Bottom) the resulting chance of picker successfully finding the key. All with $\lambda = 1000$.

of the figure shows the resulting probability that the key will be discovered using K brute-force guesses. For brevity we only included in the graph the results for when $\lambda = 1000$. Note that the differing cost models do not have much impact in this scenario and are overshadowed by the magnitude of the exponentially increasing key space that is available to the picker.

For comparison purposes, the dashed line in the figure represents the fixed picker behavior that chooses from among only the 32 bit keys and incurring a linear cost of picking, 0.5. The top of the figure shows that this strategy does worse in terms of loss than one which responds to adversary's power.

V. COSTLY-GUESSES

An alternative setting to a guesser with a limited number of guesses is one with costly actions: consider a guesser that incurs a cost of $\sigma > 0$ per each guess. To keep the model and analysis simple, we assume that this cost is not guess-dependent. In the case of passwords, for instance, this means that computation of the hash of a guess is independent of the guess itself, which largely holds for most hashing schemes.

A pure strategy of the picker and her pure strategy set are the same as in the Capped-Guesses setting: the picker selects a secret d from the set of all secrets \mathcal{P} . The guesser's strategy, however, can no longer be simply modeled as a subset of guesses to try, because here, the order of guesses matters: if the secret is found, the guesser will stop the search and save on the exploration cost. As before, the guesser is strictly better off in expectation to avoid multiple tries of the same guess. Intuitively, a pure strategy of the guesser, as his plan of action, can be represented as a sequence of guesses without repetition, i.e., a *permutation* of a subset of \mathcal{P} . The interpretation of a sample strategy $A = \langle p_1, \dots, p_\tau \rangle$ where $p_i \in \mathcal{P}$ for $1 \leq i \leq$

$\tau \leq |\mathcal{P}|$, will be the following: first try secret p_1 as a guess, if it is not correct, i.e., if the attempt fails, then try p_2 , if it fails try p_3 , and so on up to p_τ , if p_τ fails, then quit the search.

For any set \mathcal{E} , let $\text{Perm}(\mathcal{E})$ represent the set of all ordered arrangements (sequences without repetition) of all the members of \mathcal{E} , i.e., $\text{Perm}(\mathcal{E}) := \{A = \langle a_1, \dots, a_{|\mathcal{E}|} \rangle \mid \{a_1, \dots, a_{|\mathcal{E}|}\} = \mathcal{E}\}$. Moreover, let $\Psi(\mathcal{E})$ be the set of all permutations over the elements of the subsets of \mathcal{E} , i.e., $\Psi(\mathcal{E}) := \{A \mid \exists \mathcal{E}' \subseteq \mathcal{E} \text{ such that } A \in \text{Perm}(\mathcal{E}')\}$. Using these notations, we can express the strategy space of the guesser \mathcal{A} as $\mathcal{A} = \Psi(\mathcal{P})$. Note that the empty sequence, which we denote by $\langle \text{quit} \rangle$ for better presentation, is part of the strategy space of the guesser as well, representing quitting before making any guesses.

In the Appendix of the accompanied technical report[8], we show how the specification of the strategy space of the guesser in Costly-Guesses scenarios can be formally derived from the standard game theory models of sequential games with imperfect information. Specifically, it constitutes the set of *reduced* pure strategies of a guesser with *perfect recall* (he remembers his past guesses).

Given a (pure) strategy profile (d, A) , we next compute the utilities of the two players: $u_D(d, A)$ and $u_A(d, A)$. First, some notations; we extend the notion of set memberships to permutations as well, i.e., for a sequence $A = \langle a_1, \dots, a_\tau \rangle$, $d \in A$ if and only if $d \in \{a_1, \dots, a_\tau\}$. Let $\mathbf{1}_A(d)$ be the indicator function determining whether d appears on sequence A , i.e., whether $d \in A$. Let $\text{pos}_A(d)$ refer to the position of the first appearance of d on sequence A if $d \in A$, and the length of sequence A otherwise. For instance, $\text{pos}_{\langle a, b, c \rangle}(b) = 2$ and $\text{pos}_{\langle a, b, c \rangle}(e) = 3$. Then we have (compare with (1)):

$$\begin{aligned} u_D(d, A) &= -c(d) - \lambda \mathbf{1}_A(d) \\ u_A(d, A) &= \gamma \mathbf{1}_A(d) - \sigma \text{pos}_A(d) \end{aligned} \quad (8)$$

As in the Capped-Guesses setting, pure strategies may not be part of any solution concept, since a pure strategy for the picker translates to unambiguously revealing her secret. Hence we should be searching for solutions in the realm of mixed strategies. As before, let δ and α denote a mix strategy of the picker and guesser, where $\delta \in \Delta(\mathcal{P})$ and $\alpha \in \Delta(\mathcal{A})$, with the only difference that \mathcal{A} is now the set of sequences of distinct guesses, i.e., $\mathcal{A} = \Psi(\mathcal{P})$. From (8), the expected utilities of the players given a mixed strategy profile (δ, α) are:

$$\begin{aligned} u_D(\delta, \alpha) &= - \sum_{d \in \mathcal{P}} c(d) \delta(d) - \lambda \sum_{A \in \mathcal{A}} \mathbf{1}_A(d) \delta(d) \alpha(A) \\ u_A(\delta, \alpha) &= \gamma \sum_{d \in \mathcal{P}} \sum_{A \in \mathcal{A}} \mathbf{1}_A(d) \delta(d) \alpha(A) - \sigma \sum_{d \in \mathcal{P}} \sum_{A \in \mathcal{A}} \text{pos}_A(d) \delta(d) \alpha(A) \end{aligned}$$

For any $A = \langle a_i \rangle_i$, we have: $\sum_{d \in \mathcal{P}} \mathbf{1}_A(d) \delta(d) = \sum_{i=1}^{|A|} \delta(a_i)$. Moreover: $\sum_{d \in \mathcal{P}} \text{pos}_A(d) \delta(d) = \sum_{i=1}^{|A|} i \delta(a_i) + |A| (1 - \sum_{i=1}^{|A|} \delta(a_i))$. Hence:

$$u_A(\delta, A) = \gamma \sum_{i=1}^{|A|} \delta(a_i) - \sigma \left[\sum_{i=1}^{|A|} i \delta(a_i) + |A| (1 - \sum_{i=1}^{|A|} \delta(a_i)) \right] \quad (9)$$

Game 2: Costly-Guesses

Players: PICKER, GUESSER

Strategy Sets: PICKER'S: $\{d \in \mathcal{P}\}$

GUESSER'S: $\{A | \exists \mathcal{E} \subseteq \mathcal{P} \text{ such that } A \in \text{Perm}(\mathcal{E})\}$

Utilities: PICKER: $u_D(d, A) = -c(d) - \lambda \mathbf{1}_A(d)$,

GUESSER: $u_A(d, A) = \gamma \mathbf{1}_A(d) - \sigma \text{pos}_A(d)$

An alternative method to derive the expression for $u_A(\delta, A)$ is the following: $\sum_{i=1}^{|A|} \delta(a_i)$ is just the probability that any of the tries on sequence A is the correct guess. Given δ and A , the search reaches a_i in A with probability $1 - \sum_{j=1}^{i-1} \delta(a_j)$. Hence, the expected number of tries is $\sum_{i=1}^{|A|} (1 - \sum_{j=1}^{i-1} \delta(a_j))$. Therefore:

$$u_A(\delta, A) = \gamma \sum_{i=1}^{|A|} \delta(a_i) - \sigma \sum_{i=1}^{|A|} \left[1 - \sum_{j=1}^{i-1} \delta(a_j) \right] \quad (10)$$

This is equivalent to the expression in (9). In our analysis, we will use either one of the two forms based on convenience.

All of the solution concepts introduced in Section III can be identically defined here as well. We will explore them in detail in the next section.

VI. ANALYSIS OF THE COSTLY-GUESSES SCENARIO

Before we delve into the analysis of the Costly-Guesses scenario, we present a simple yet instrumental lemma:

Lemma 2: Let \mathcal{E} be a non-empty subset of \mathcal{P} , and let $\text{unif}(\mathcal{E})$ represent the uniform distribution over \mathcal{E} , i.e., $\delta = \text{unif}(\mathcal{E})$ if and only if $\delta(p) = \mathbf{1}_{\mathcal{E}}(p)/|\mathcal{E}|$. Then, for any $A \in \text{Perm}(\mathcal{E})$, $u_A(\text{unif}(\mathcal{E}), A) = \gamma - (|\mathcal{E}| + 1)\sigma/2$, i.e., the expected utility of the guesser for any strategy that exhausts \mathcal{E} is $\gamma - (|\mathcal{E}| + 1)\sigma/2$.

Proof: The secret is a member of \mathcal{E} , hence it will be found with certainty, yielding the positive gain of γ . Each guess costs the guesser σ . The number of guesses before (and including) the correct one is i with probability $1/|\mathcal{E}|$. Hence the expected number of tries is $\sum_{i=1}^{|\mathcal{E}|} i/|\mathcal{E}| = (|\mathcal{E}| + 1)/2$. ■

We will investigate the maximin and minimax strategies of the picker first. The picker's maximin strategy is choosing a secret from the cheapest partition, i.e., a picking strategy $\delta \in \Delta(\mathcal{P})$ is maximin if and only if $\sum_{p \in \mathcal{E}_1} \delta(p) = 1$. To see this, note that a strategy of the guesser that explores all of the possible secrets, i.e., a permutation of the entire \mathcal{P} , minimizes the utility of the picker irrespective of the choice of her strategy. Hence, facing this worst case strategy of the guesser, the picker must only select from the cheapest partition.

A minimax strategy of the picker, on the other hand, is uniform randomization over the entire \mathcal{P} , due to the following two intuitive lemmas:

Lemma 3: Let \mathcal{E} be a non-empty subset of \mathcal{P} . Then, for any $\delta \in \Delta(\mathcal{P})$ such that $\text{supp}(\delta) \subseteq \mathcal{E}$, we have: $\sup_{\alpha \in \Delta(A)} u_A(\delta, \alpha) \geq \sup_{\alpha \in \Delta(A)} u_A(\text{unif}(\mathcal{E}), \alpha)$.

Lemma 4: Let $\mathcal{E}, \mathcal{E}'$ be two non-empty subsets of \mathcal{P} such that $|\mathcal{E}| \leq |\mathcal{E}'|$. Then, $\sup_{\alpha \in \Delta(A)} u_A(\text{unif}(\mathcal{E}), \alpha) \geq \sup_{\alpha \in \Delta(A)} u_A(\text{unif}(\mathcal{E}'), \alpha)$.

The first lemma simply confirms that uniform distribution gives the least amount of useful information to the guesser. The second lemma states that uniform randomization over a bigger set is guaranteed not to help the guesser. Proof of Lemma 3 is in the technical report[8]. Lemma 4 follows directly from Lemma 2.

As we can see, in the Costly-Guesses setting, the strategically pessimistic and the sheer antagonistic plans of action for the picker (her maximin and maximin strategies, respectively) lead to uninteresting extremes, suggesting that rationality consideration of both players have a more decisive role. Next, we turn our attention to NE solutions.

A. Costly-Guesses: Nash Equilibria

When $\gamma < \sigma$, the cost of trying even a single guess exceeds the gain of finding the secret. Hence, irrespective of the strategy of the picker, the guesser never enters the game:

Proposition 2: In a Costly-Guesses game, if $\gamma < \sigma$, then in all NE (δ^*, α^*) , we have: $\delta^*(\mathcal{E}_1) = 1$ and $\alpha^*(\{\text{quit}\}) = 1$, i.e., the picker chooses from the cheapest partition and the guesser does not make any attempt.

What happens when $\gamma > \sigma$? If $\gamma < (1 + \sum_{i=1}^M |\mathcal{E}_i|)\sigma/2$ for some $M \leq N$, then following Lemma 2, the picker can dissuade the guesser from entering the game by uniformly randomizing over the first M partitions. When the picker assumes a high cost for losing her secret, i.e., for large values of λ , this seems to be something she will opt for. However, our next proposition reveals that, surprisingly, if there is no partition that is big enough that uniform randomization over it *alone*, i.e., single-handedly, can dissuade the guesser from entering, then in *all* NE of the game, the picker chooses a cheapest secret and loses it with certainty, and remarkably, this is true *irrespective* of the magnitude of λ :

Proposition 3: In a Costly-Guesses game, if $\gamma > (1 + |\mathcal{E}_i|)\sigma/2$ for all i for which $C_i < C_1 + \lambda$, then in all NE (δ^*, α^*) , we have: $\delta^*(\mathcal{E}_1) = 1$ and $u_D(\delta^*, \alpha^*) = -C_1 - \lambda$, i.e., the picker chooses only from the cheapest partition and the guesser finds it with certainty.

The detailed proof of the proposition is provided in the technical report[8]. Here we provide an informal summary of the proof with the aim of giving an idea why we have this "failure" of NE for the picker: in any NE, the mixed strategies of the two players must be best responses to each other. Therefore, in a NE, the picker only assigns positive probability of selection from costlier partitions because of the threat imposed by the exploration probabilities of the guesser. Suppose there is a NE in which the picker assigns strictly positive probabilities to secrets from partitions \mathcal{E}_1 to \mathcal{E}_M . This means that the guesser explores \mathcal{E}_1 to \mathcal{E}_M with strictly decreasing probabilities. This in turn implies that the guesser must find it a best response to explore \mathcal{E}_{M-1} and not \mathcal{E}_M among his set of best responses that he randomizes over. Note that the picker never assigns a strictly higher probability to

members from a costlier partition. This means that if exploring \mathcal{E}_{M-1} and not \mathcal{E}_M must be a best response of the guesser, so must be exploring \mathcal{E}_1 through \mathcal{E}_{M-1} and not \mathcal{E}_M . However, this can never be the case: if the guesser explores all of the partitions \mathcal{E}_1 to \mathcal{E}_{M-1} and fails, then given the randomization of the picker, he is now certain that the secret is in \mathcal{E}_M . Given the condition $\gamma > (\mathcal{E}_M + 1)\sigma/2$, the guesser is strictly better off to continue to explore \mathcal{E}_M as well. Hence, the starting assumption about the NE strategy of the picker could not be true.

The next proposition shows what may happen when the condition of Prop. 3 is relaxed (proof is in the technical report):

Proposition 4: In a Costly-Guesses game where $\gamma > \sigma$, if $\exists M = \min\{i | \gamma < (|\mathcal{E}_i| + 1)\sigma/2, C_i \leq C_1 + \lambda\}$, then in all NE (δ^*, α^*) we have: $u_D(\delta^*, \alpha^*) \leq -C_M$.

This proposition does not quite redeem the stark situation with NE solutions for the picker. For instance, consider a case where the picker could prevent the guesser from entering the game by randomizing over \mathcal{E}_1 and \mathcal{E}_2 , and the cheapest partition that is big enough to single-handedly prevent the guesser from entering the game is \mathcal{E}_3 . Then the picker has to settle for a cost of C_3 , which can be much larger than any weighted average of C_1 and C_2 . Moreover, the proposition only provides a (tight) upper-bound on the expected utility of the picker among all NE. That is, $-C_M$ is the expected utility of the picker in the best NE for her, and worse NE for the picker can still exist. In particular, if $\gamma > (|\mathcal{E}_1| + 1)\sigma/2$, then (δ^*, α^*) where $\delta^* = \text{unif}(\mathcal{E}_1)$ and $\alpha^* = \text{unif}(\text{Perm}(\mathcal{E}_1), \text{Perm}(\cup_{i=2}^N \mathcal{E}_i))$ is also technically a NE: given that the picker chooses uniformly from the cheapest partition, it is a best response for the guesser to explore the whole set of secrets starting from the cheapest partition; likewise if the guesser's strategy is to explore the whole set of secrets, then the guesser's best response is to choose from the cheapest partition, since she will lose her secret anyway. This NE, as in Prop. 3, yields for the picker the worst possible in any NE: her maximin utility, that is $-C_1 - \lambda$.

What causes the poor performance of the picker in NE is the absence of a credible commitment to a deterring randomization. Indeed the picker prefers to induce the guesser to abstain, however, if the guesser is not going to enter the game, the picker prefers to select a least costly secret. The picker can remove this possibility from the reasoning of the guesser by credibly communicating a commitment to a mixed strategy. This is exactly the setup for Strong Stackelberg Equilibria, which we analyze next.

B. Equilibrium Example: Passwords

Figure 5 shows the result of equilibrium behavior on the loss of the picker in the costly guesses model for password selection. Most of the figure is a lower bound on loss as per Prop 2. For low ratios of γ/σ , the guesser does not participate at all and results only in the cost of picking the simplest password. For large enough ratios, the picker gives up, incurring a loss of λ and the cost of the simplest password.




Fig. 5. RockYou-based password picker loss in Nash equilibrium or lower bound (gray) and in Stackelberg equilibrium (black dotted) as function of γ/σ . For all, $\lambda = 2$.




Fig. 6. Picker loss in the Nash (solid lines) and Stackelberg (dotted lines) key picking strategy as a function of γ/σ , for linear and cubic cost functions (lighter line is cubic cost). For all, $\lambda = 2$.

In the mid-range cases, the loss factor λ plays no role.

C. Equilibrium Example: Cryptographic Keys

The solid lines in Figure 6 show picker loss for the key selection scenario with costly guesses. The results contain essentially the same features as the password selection example: low-enough ratios of γ/σ results in guesser not participating, high enough ratios result in the picker giving up (this is out of the frame at picker loss equal to $\lambda = 2$), and in between the picker exhibits an increasing expected loss.

D. Costly-Guesses: Strong Stackelberg Equilibria

Here we assume the picker has access to an apparatus that enables her to credibly communicate a commitment to a mixed strategy to the guesser. We develop the optimal randomizations for the picker given the fact that the guesser, observing the committed randomization, best-responds to it. Formally, we derive the SSE strategies of the picker.

First, note that if $\gamma < \sigma$, then irrespective of the choice of the picker, the guesser will never attempt a guess. Then the SSE strategy of the picker for these cases is, trivially, a choice from the cheapest partition, yielding the picker a utility of $-C_1$ and the guesser, zero. Therefore, in the rest of this section, we only consider $\gamma > \sigma$. We show the following: if at all worth protecting the secret, the picker should commit to a randomization that makes not entering the game a best response for the guesser, i.e., the cheapest randomization that leaves the guesser indifferent between entering the game and quitting at the beginning. In particular, committing to randomizations that leave incentive for the guesser to perform

even a partial search is never optimal.⁸ We specifically develop a linear optimization that gives the SSE strategy of the picker.

Proposition 5: Consider the following linear programming:

$$u_D^* = \max_{\nu_i} \left[- \sum_{i=1}^N C_i \nu_i \right] \quad \text{subject to :}$$

$$\nu_i \geq 0 \text{ for } 1 \leq i \leq N, \quad \sum_{i=1}^N \nu_i = 1, \quad \frac{\nu_i}{|\mathcal{E}_i|} \geq \frac{\nu_{i+1}}{|\mathcal{E}_{i+1}|} \text{ for } 1 \leq i \leq N-1$$

$$\gamma \sum_{i=1}^K \nu_i - \sigma \sum_{i=1}^K \left[|\mathcal{E}_i| \left(1 - \sum_{j=1}^{i-1} \nu_j \right) - \frac{|\mathcal{E}_i| - 1}{2} \nu_i \right] \leq 0 \text{ for } 1 \leq K \leq N$$

For $(|\mathcal{P}| + 1)\sigma/2 > \gamma$, the LP is feasible. Let $(\nu_1^*, \dots, \nu_N^*)$ be a solution. If $u_D^* > -C_1 - \lambda$, then a SSE strategy of the picker is $\delta(p) = \nu_i^*/|\mathcal{E}_i|$ for $p \in \mathcal{E}_i$. If $u_D^* < -C_1 - \lambda$, the SSE strategy of the picker is to simply choose a secret from the cheapest partition (which induces the guesser to enter, explore that partition and find the secret with certainty). Same is true when $(|\mathcal{P}| + 1)\sigma/2 < \gamma$.⁹

The proof of the proposition is provided in the technical report. Note that when $(|\mathcal{P}| + 1)\sigma/2 < \gamma$, following Lemma 2, even uniform randomization over the entire set of \mathcal{P} does not deter the guesser from entering the game and exploring the whole secret space, as it yields him a strictly positive utility of $\gamma - (|\mathcal{P}| + 1)\sigma/2$. Since uniform randomization is a minimax strategy of the picker (intuitively, it gives the least useful information to the guesser), any other randomization also results in a strictly positive utility for full exploration of the guesser. This means the best strategy of the picker is then choosing from a cheapest partition, since she will lose her secret to the guesser anyway.

When $(|\mathcal{P}| + 1)\sigma/2 > \gamma$, uniform randomization over a subset of secrets can lead to a negative expected utility of the guesser for entering the game and exploring any portion of the secret space. However, our numerical examples of the proposition reveal that the cheapest randomization that achieves this goal is almost never completely uniform (or even necessarily uniform over the union of some cheapest partitions except for the costliest of them).

E. Stackelberg Examples

The difference between the Nash equilibrium and the Stackelberg equilibrium is demonstrated in Figure 5 for the password picking example and in Figure 6 for the key selection example. In both, the picker’s loss in Stackelberg equilibrium as a function of γ/σ is denoted by dotted lines. In the case of key selection, linear and cubic cost models are shown with linear as dark lines and cubic as light lines. The Stackelberg strategies can be seen to perform better than the Nash strategies shown as solid lines.

⁸This is reminiscent of this pithy quote [20] from Zhuge Liang, a recognized ancient Chinese military strategist and statesman: “*The wise win before they fight, while the ignorant fight to win.*”

⁹One can find uniqueness conditions for the SSE strategy of the picker, using standard results in linear programming (e.g. [21]). However, the uniqueness of the utility of the picker follows from the optimization itself.




Fig. 7. The distribution of key sizes in the Stackelberg key picking strategy for a variety of γ/σ values.

Figure 7 demonstrates the Stackelberg strategy for key selection in more detail for three different γ/σ values. The selection of key in these solutions is mostly uniform among keys up to a certain length except for larger keys whose probability sharply falls off. The range of key sizes which are selected over increases with γ/σ .

VII. DISCUSSIONS

Policy implications of picker’s optimal strategies in Capped-Guesses: System administrators usually use password selection rules (composition policies) to increase the entropy of the passwords selected by users. The choice of an optimal rule-set has been a topic of research [10, 13, 14, 22]. Recall that an interpretation of the partitions of the secret space based on their usability costs was that these partitions can be assumed to satisfy increasingly complex password composition rules. Hence, our Prop. 1 suggests that the search for the “optimal” composition rule is amiss. The optimal secret picking strategy is a uniform randomization over a “union” of these partitions, and hence, no single composition rule is optimal. Our result suggests that optimal composition rules are generated by randomizing across different rule-set, and specifically, each composition rule should be prompted to a user with a probability that is proportional to the size of password space created by that rule.

Credible commitment to a randomization Recall that in a Capped-Guesses scenario, the SSE, NE and Maximin strategies of the picker turn out to be identical. In particular, there is no gain in communicating a commitment to the adversaries. In the Costly-Guesses model, however, a credible commitment to a randomization makes a substantial difference, and is critical to prevent the failure of NE. Hence, in such situations, it is not sufficient to have access to a randomization device, but further, the randomization should be made public knowledge and verifiable to become credible.

Optimal attacks in Capped-Guesses Our Prop. 1 suggest that in a Capped-Guesses attack, e.g. using pre-computed tables, even facing a rational defender that plays optimally and hence uses uniform randomization, the adversary must choose passwords randomly from the whole selection range of the user, however, should choose simpler passwords with more probability and include more difficult ones with increasingly less probabilities.

Interpretation of mixed strategies The game theoretic solutions that we developed involved randomization. Specifically, in mixed NE, each player’s randomization leaves the other indifferent across his/her randomization support. Although these behaviors can be explicitly associated with deliberate randomization or through the use of randomization devices (e.g. when a random key generator algorithm is used), these are not the only way such equilibria can be interpreted. Without going to the details [23], we just mention some of the alternative interpretations equilibrium solution involving mixed strategies. Namely, the probabilities can represent (a) time averages of player’s behavior that exploit an “adaptive” process, (b) fractions of the total “population” of each player that adopt pure strategies, (c) limits of pure strategy Bayesian equilibria where each player is slightly uncertain about the payoffs of the others, and (d) A “consistent” set of “beliefs” that each player has about the other regarding their behavior.

Other applications: Finally, it is worth mentioning that even though we motivated our models based on password and cryptographic key selection, the generality of the model allows it to be applicable to other contexts as well. As an example of a completely different context but with identical abstraction, consider a user that aims to send a convoy from a source to a destination over a transport network, or transmit a packet over a communication network. There are multiple paths available and the user’s objective is to use this path diversity to minimize the risk of being intercepted on a path by an adversary. However, the paths may have different utilities as some may provide lower delays and higher quality of service, a preference that can be exploited by adversaries as well.

VIII. RELATED WORK

User password selection and attacks has been extensively studied in the literature [2, 10, 13–15, 22], and due to its practical significance, continues to be a hot area of research [3]. These works generally aim at evaluating the efficacy of password attacks as well as measuring the strength of different password composition rules through statistical metrics. In contrast to our work, these papers consider the user or the adversaries one at a time, as opposed to considering that both parties will adapt to each other’s choice of policies. Analysis of such strategic actions and reactions can be done through a game theoretic framework, which to our best of knowledge, our work is the first in this context.

Game and decision theory has been applied in other cybersecurity contexts with promising potentials [24, 25]. The first part of our work (Capped-Guesses) is, in its abstract form, similar to the security game model analyzed in [26]. In their model, the defender has limited resources to cover a wide range of targets, while an adversary chooses a single target to attack. If targets are thought of as secrets, the defender in their model is akin to the guesser in our work, and their adversary is our picker. Therefore, our Capped-Guesses model is the “complement” of their model. Specifically, the results that they develop for their defender will be translatable to our guesser. However, the focus of our paper was on the picker.

Another line of research from theoretical game theory is search theory and search games [27]. Existence of user preferences over the secrets to pick from is missing from such models. However, such preferences are at the heart of usability-security trade-off settings investigated in our paper.

IX. CONCLUSION

We developed tractable game-theoretic models that capture the essence of secret picking vs guessing attacks in the presence of preferences over the secret space. We then provided a full analysis of our models with the aim of investigating fundamental trends and properties in the design of secret-picking policies that attain optimal trade-offs between usability and security, taking into account the exploitation of the knowledge of such trade-offs by an adversary. Notably, we computed the secret picking policies that are optimal with respect to a range of strategic metrics (Maximin, Minimax, Nash Equilibria, Stackelberg Equilibria).

We distinguished between two classes of guessing attacks: those in which the number of available guesses to an adversary is capped (Capped-Guesses), and those in which an adversary has potentially unlimited number of tries but incurs a cost per each guess (Costly-Guesses). Our analysis revealed the crucial role that such distinction between the nature of the guessing adversary plays on the expected outcome. Specifically, we showed that in the Capped-Guesses settings, the NE strategy of the secret picker is still uniform but over a low-cost subset of the secret space, where the size of the subset depends on the parameters of the adversary only through the number of available guesses. In contrast, we established that for Costly-Guesses scenarios, except for trivial cases, NE fails to attain a desirable outcome for the secret picker. For this setting, we showed how deterrence of adversaries as her optimal strategy crucially depend on existence of a credible commitment to a randomization strategy. We illustrated our results through a series of numerical examples using real-world data-sets.

Future Directions: One of the main areas of extending this work is dealing with uncertainty in the parameters of the players. For instance, the picker may not accurately know the type of the guesser or their guessing size cap or their guessing costs. One approach to formally take such uncertainties into account is a Bayesian game approach, for which, this work lays the foundation of.

Moreover, in this paper, we assumed that once the secret is selected, the picker does not get to change it later, either as a blind (open-loop) policy or as a reaction to some signal generated by the actions of an adversary. Note that if the act of changing the secret does not bring any cost to the picker, and both parties are aware of secret-changing occasions, then our results are still applicable, since in essence, the two players play the same game after each reset. However, the previous choices of the picker may affect her future utilities, and hence the whole game. For instance, the act of changing the secret may be costly for the picker, or changing the secret only slightly may be associated with less cost than changing it drastically. In such scenarios, a rational adversary can exploit

such preferences and carry some useful information from each round of the game to boost his overall attack. Investigation of such scenarios using dynamic game theory is a potential extension of our work.

Another interesting scenario to investigate is when the picker is choosing multiple secrets, where there is an increasing loss for the number of secrets guessed correctly by an adversary. The two extremes are (1) when the guesser wins if any of the secrets are discovered, and (2) when the guesser wins only if all of the secrets are discovered.

X. ACKNOWLEDGMENTS

This research was partially sponsored by US Army Research Laboratory and the UK Ministry of Defence under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defence, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. The first and third authors were also supported by the Project “Games and Abstraction: The Science of Cyber Security” funded by EPSRC, Grants: EP/K005820/1, EP/K006010/1.

REFERENCES

- [1] J. Bonneau, C. Herley, P. C. Van Oorschot, and F. Stajano, “The quest to replace passwords: A framework for comparative evaluation of web authentication schemes,” in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 553–567.
- [2] J. Bonneau, “Guessing human-chosen secrets,” Ph.D. dissertation, University of Cambridge, 2012.
- [3] D. Florêncio, C. Herley, and P. C. Van Oorschot, “An administrator’s guide to internet password research,” in *Proceedings of the 28th USENIX conference on Large Installation System Administration*. USENIX Association, 2014, pp. 35–52.
- [4] Verizon Business Risk Team, “Data breach investigations report,” *Studie, erhältlich unter www.verizonbusiness.com*, 2008.
- [5] M. A. Sasse, S. Brostoff, and D. Weirich, “Transforming the ‘weakest link’—a human/computer interaction approach to usable and effective security,” *BT technology journal*, vol. 19, no. 3, pp. 122–131, 2001.
- [6] J. Yan *et al.*, “Password memorability and security: Empirical results,” *IEEE Security & privacy*, no. 5, pp. 25–31, 2004.
- [7] X. Boyen, “Halting password puzzles,” in *Proc. Usenix Security*, 2007.
- [8] M. Khouzani, P. Mardziel, C. Cid, and M. Srivatsa, “Picking vs. guessing secrets – a game-theoretic analysis (technical report),” July 2015, available online at: http://www.eecs.qmul.ac.uk/~khouzani/Reports/CSF15_TechRep.pdf.
- [9] H. W. Kuhn, “Introduction to john von neuman and oskar morgenstern’s theory of games and economic behavior,” *Introductory Chapters*, 2007.
- [10] J. Bonneau, “The science of guessing: analyzing an anonymized corpus of 70 million passwords,” in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 538–552.
- [11] M. Zviran and W. J. Haga, “Password security: an empirical study,” *Journal of Management Information Systems*, pp. 161–185, 1999.
- [12] A. Vance, “If your password is 123456, just make it hackme,” <http://www.nytimes.com/2010/01/21/>, 2010.
- [13] M. Weir, S. Aggarwal, M. Collins, and H. Stern, “Testing metrics for password creation policies by attacking large sets of revealed passwords,” in *Proceedings of the 17th ACM conference on Computer and communications security*. ACM, 2010, pp. 162–175.
- [14] P. G. Kelley, S. Komanduri, M. L. Mazurek, R. Shay, T. Vidas, L. Bauer, N. Christin, L. F. Cranor, and J. Lopez, “Guess again (and again and again): Measuring password strength by simulating password-cracking algorithms,” in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 523–537.
- [15] B. Ur, P. G. Kelley, S. Komanduri, J. Lee, M. Maass, M. L. Mazurek, T. Passaro, R. Shay, T. Vidas, L. Bauer *et al.*, “How does your password measure up? the effect of strength meters on password creation.” in *USENIX Security Symposium*, 2012, pp. 65–80.
- [16] J. Katz and Y. Lindell, *Introduction to modern cryptography*. CRC Press, 2014.
- [17] T. Basar, G. J. Olsder, G. Clsder, T. Basar, T. Baser, and G. J. Olsder, *Dynamic noncooperative game theory*. SIAM, 1995, vol. 200.
- [18] B. Von Stengel and S. Zamir, “Leadership games with convex strategy sets,” *Games and Economic Behavior*, vol. 69, no. 2, pp. 446–457, 2010.
- [19] J. Letchford, D. Korzhyk, and V. Conitzer, “On the value of commitment,” *Autonomous Agents and Multi-Agent Systems*, pp. 1–31, 2012.
- [20] D. McAdams, *Game-Changer: game theory and the art of transforming strategic situations*. WW Norton & Company, 2014.
- [21] O. L. Mangasarian, “Uniqueness of solution in linear programming,” *Linear algebra and its applications*, vol. 25, pp. 151–162, 1979.
- [22] S. Komanduri, R. Shay, P. G. Kelley, M. L. Mazurek, L. Bauer, N. Christin, L. F. Cranor, and S. Egelman, “Of passwords and people: measuring the effect of password-composition policies,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2011, pp. 2595–2604.
- [23] H. Gintis, *Game theory evolving: A problem-centered introduction to modeling strategic behavior*. Princeton University Press, 2000.
- [24] M. Tambe, *Security and game theory: Algorithms, deployed systems, lessons learned*. Cambridge University

Press, 2011.

- [25] T. Alpcan and T. Başar, *Network security: A decision and game-theoretic approach*. Cambridge University Press, 2010.
- [26] Z. Yin, D. Korzhyk, C. Kiekintveld, V. Conitzer, and M. Tambe, “Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness,” in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 1139–1146.
- [27] S. Alpern, R. Fokkink, G. Leszek, R. Lindelauf, V. Subrahmanian *et al.*, *Search theory: a game theoretic perspective*. Springer Science & Business Media, 2013.