

A birthday present every eleven wallets? The security of customer-chosen banking PINs

Joseph Bonneau, Sören Preibusch, Ross Anderson

Computer Laboratory
University of Cambridge
{jcb82,sdp36,rja14}@cl.cam.ac.uk

Abstract. We provide the first published estimates of the difficulty of guessing a human-chosen 4-digit PIN. We begin with two large sets of 4-digit sequences chosen outside banking for online passwords and smartphone unlock-codes. We use a regression model to identify a small number of dominant factors influencing user choice. Using this model and a survey of over 1,100 banking customers, we estimate the distribution of banking PINs as well as the frequency of security-relevant behaviour such as sharing and reusing PINs. We find that guessing PINs based on the victims' birthday, which nearly all users carry documentation of, will enable a competent thief to gain use of an ATM card once for every 11–18 stolen wallets, depending on whether banks prohibit weak PINs such as 1234. The lesson for cardholders is to never use one's date of birth as a PIN. The lesson for card-issuing banks is to implement a denied PIN list, which several large banks still fail to do. However, blacklists cannot effectively mitigate guessing given a known birth date, suggesting banks should move away from customer-chosen banking PINs in the long term.

1 Introduction

Personal Identification Numbers, or PINs, authenticate trillions of pounds in payment card transactions annually and are entrenched by billions of pounds worth of infrastructure and decades of customer experience. In addition to their banking role, 4-digit PINs have proliferated in a variety of other security applications where the lack of a full keypad prevents the use of textual passwords such as electronic door locks, smartphone unlock codes and voice mail access codes. In this work, we provide the first extensive investigation of the security implications of human selection and management of PINs.

1.1 History of PINs

We refer the reader to [4] for a good overview of the history of banking cards and ATMs; we summarise the development of PINs for security here. The historical record suggests that PINs trace their origins to automated dispensing and control systems at petrol filling stations. In the context of banking, PINs first appeared in separate British cash machines deployed in 1967, with 6-digit PINs in the

Barclays-De La Rue system rolled out in June and 4-digit PINs in the National-Chubb system in September. According to John Shepherd-Barron, leader of the De La Rue engineering team, after his wife was unable to remember six random digits he reduced the length to four.

Early cash machines were stand-alone, offline machines which could only exchange cash for punched cards (which were kept by the machine). The primary use case was to cease branch operations on Saturdays and still allow customers to retrieve cash. Interestingly, cash machines deployed contemporaneously in Japan and Sweden in 1967 used no PINs and absorbed losses from lost or stolen cards. As late as 1977, Spain's La Caixa issued cards without PINs.

PINs were initially bank-assigned by necessity as they were hard-coded onto cards using steganographic schemes such as dots of carbon-14. Soon a variety of schemes for storing a cryptographic transformation of the PIN developed.¹ The IBM 3624 ATM controller introduced an influential scheme for deriving PINs in 1977 [5]. PIN verification consisted of a DES encryption of the user's account number, converting the first 4 hexadecimal digits of the result into decimal using a lookup table, adding a 4-digit PIN offset modulo 10^4 , and comparing to the entered PIN. Changing the PIN offset stored on the card enabled the user to choose their own PIN. Banks began allowing customer-chosen PINs in the 1980s as a marketing tactic, though it required substantial infrastructure changes.

The development of Visa and MasterCard and the interconnection of ATM networks globally in the 1990s cemented the use of PINs for payment card authentication in both the 1993 ISO 9564 standard [3] and 1995 EMV standard [1]. Today, most cards use the Visa PVV scheme, which stores a DES-based MAC of the account number and PIN called the pin-verification value (PVV) which can be re-computed to check if a trial PIN is correct.

The EMV standard further led to PINs taking on the role of authorising payments at merchant tills, with the card's chip verifying the customer's PIN internally.² Technically, this use of PINs uses a different mechanism than that for ATM authentication, though in all practical deployments the two PINs are the same and may only be changed at an ATM. With the advent of EMV, PINs must be entered more often and into a plethora of vendor terminals, increasing the risk of compromise.

Chip cards have also enabled the deployment of hand-held Chip Authentication Program (CAP) readers since 2008 for verifying Internet transactions [10]. CAP readers allow muggers to verify a PIN demanded from a victim during an attack; they can also be used to guess offline the PIN on a found or stolen card.

1.2 Standards and practices in PIN selection

Published standards on PIN security provide very brief treatment of human factors. The EMV standard [1] requires support for PINs of 4–12 digits, in line

¹ James Goodfellow patented a cryptographic PIN derivation scheme in 1966 [12]. Amongst others, he has been called the inventor of PINs and ATMs.

² EMV was deployed in the UK from 2003 under the branding "Chip and PIN." It is now deployed in most of Europe, though notably not in the United States.

with earlier Visa standards, but makes no mention of PIN selection. Separately, Visa maintains *Issuer PIN Security Guidelines* with several recommendations for users, specifically that they never write down their PIN or use it for any other purpose. The document is neutral between issuer-assigned PINs or customer-chosen PINs, providing one sentence about PIN selection [2]: “Select a PIN that cannot be easily guessed (i.e., do not use birth date, partial account numbers, sequential numbers like 1234, or repeated values such as 1111).”

ISO 9564 [3] covers PIN security and is largely similar to Visa’s guidelines, mostly focusing on PIN transmission and storage. It adds a recommendation against “historically significant dates,” and PINs chosen as words on the keypad. Neither standard mentions using a “denied PIN list” to blacklist weak PINs, as is recommended in standards for text passwords [8].

As a result of the vague standards, PIN requirements vary significantly but the minimal 4-digit length predominates. PIN length appears integrated into cultural norms: there is rarely variation within competitive regions, while in some locales most card issuers require PINs longer than 4 digits.³ Similarly, most banks allow user-chosen PINs, with a few regional exceptions such as Germany.

Because denied PIN lists aren’t publicly advertised, we evaluated several banking cards by requesting the PIN 1234.⁴ In the UK, this was denied by Barclays, HSBC and NatWest but allowed by Lloyds TSB and The Co-op Bank. In the USA, this was denied by Citibank and allowed by Bank of America, HSBC and Wells Fargo. We only identified card-specific denied PIN lists; we found no ATM implementing local restrictions. At one bank we tested, Chase in the USA, self-service PIN changes are not possible and changes must be made in-person. Banks’s policies may vary with time or location (note the inconsistency of HSBC between the USA and UK), but denied PIN lists are clearly not universal.

1.3 Academic research

Research on authentication systems involving human-chosen secrets consistently finds that people favour a small number of popular (and predictable) choices. Strong bias has been analysed for textual passwords starting with Morris and Thompson in 1979 [15] and confirmed in many studies since [19]. Similar bias has been identified in responses to personal knowledge questions [6] and in graphical password schemes [20]. Despite their wide deployment, there exists no academic research about human selection of PINs.

The best-known research on PINs, such as Murdoch et al.’s “no-PIN attack” [16], has identified technical flaws in the handling and verification of PINs but not addressed PIN guessing. Kuhn identified in 1997 that the use of unbalanced decimalisation tables introduced a bias into the distribution of PIN offsets which could be exploited by an attacker to improve PIN guessing [13]. Bond and Zieliński developed further decimalisation-based attacks in 2003 [5]. Both attacks can be improved with knowledge of human tendencies in PIN selection.

³ For example, banks in Switzerland assign 6–8 digit PINs, and banks in Italy typically use 5-digit PINs. Canadian banks use a mix of 4-digit and 6-digit PINs.

⁴ We assume any reasonable denied PIN list would include 1234 and allowing this PIN indicates no restrictions exist.

2 Quantifying resistance to guessing

We consider abstractly the probability distribution⁵ of PINs \mathcal{X} over the set $\{0000, \dots, 9999\}$. We consider each PIN x_i to have probability p_i , with $p_1 \geq p_2 \geq \dots \geq p_N$. Several works have formally treated the mathematics of guessing an unknown value $X \stackrel{R}{\leftarrow} \mathcal{X}$ [14,9,7,17]. We use the notation and terminology from [6] throughout this paper.

A traditional measure of guessing difficulty is Shannon entropy:

$$H_1 = - \sum_{i=1}^N (p_i \cdot \log_2 p_i) \quad (1)$$

However, this is mathematically unsuited to measuring guessing difficulty⁶ [14,9] and recently been confirmed experimentally to be a poor measure of cracking difficulty for human-chosen passwords [21]. A more sound measure is *guesswork*:

$$G = \sum_{i=1}^N (p_i \cdot i) \quad (2)$$

G represents the expected number of sequential guesses to determine X if an attacker proceeds in optimal order [14]. Both G and H_1 are influenced by rare events significantly enough to make them misleading for security analysis. A preferable alternative is *marginal guesswork* μ_α , which measures the expected number of guesses required to succeed with probability α :

$$\mu_\alpha = \min \left\{ j \in [1, N] \mid \sum_{i=1}^j p_i \geq \alpha \right\} \quad (3)$$

In particular, $\mu_{0.5}$, representing the number of attempts needed to have a $1/2$ chance of guessing correctly, has been suggested as a general alternative to G , as it is less influenced by low-probability events [17]. In the case of PINs, attackers are almost always externally limited in the number of guesses they can try. In this case, the best metric is the *marginal success rate* λ_β , the probability that an attacker can correctly guess X given β attempts:

$$\lambda_\beta = \sum_{i=1}^\beta (p_i) \quad (4)$$

Locking a payment card after 3 incorrect guesses is standard practice. However, different counters are used for ATM requests and payment requests, meaning a thief with a CAP reader and access to an ATM can typically make 6 guesses. Thus, we are primarily concerned with estimating λ_3 and λ_6 , though other values are of interest if a user has reused a PIN for multiple cards.

⁵ The distribution may vary between different populations, or with knowledge of auxiliary information (such as a card holder's birthday).

⁶ Shannon entropy represents the average number of bits needed to encode a variable $X \stackrel{R}{\leftarrow} \mathcal{X}$. It measures the expected number of yes/no queries an attacker must make about the membership of X in subsets $\mathcal{X}' \subset \mathcal{X}$, which is fundamentally different from guessing individual values.

distribution	H_1	\tilde{G}	$\tilde{\mu}_{0.5}$	λ_3	λ_6
RockYou 4-digit sequences	10.74	11.50	9.11	8.04%	12.29%
RockYou regression model	11.01	11.79	9.39	5.06%	7.24%
iPhone unlock codes	11.42	11.83	10.37	9.23%	12.39%
iPhone regression model	11.70	12.06	10.73	9.21%	11.74%
random 4-digit PIN	13.29	13.29	13.29	0.03%	0.06%

Table 1. Guessing metrics for 4-digit sequences in RockYou passwords and iPhone unlock codes. Values are also shown for the regression-model approximation for each distribution and for a uniform distribution of 4-digit PINs.

The metrics are not directly comparable, as H_1 is in units of bits, G and μ_α in units of guesses, and λ_β is a probability. It can be helpful to convert all of the metrics into bits by taking the base-2 logarithm of a uniform distribution which would have the same value of the metric, as demonstrated in [6]:

$$\tilde{G} = \log_2(2 \cdot G(\mathcal{X}) - 1); \quad \tilde{\lambda}_\beta = \log_2\left(\frac{\beta}{\lambda_\beta(\mathcal{X})}\right); \quad \tilde{\mu}_\alpha = \log_2\left(\frac{\mu_\alpha(\mathcal{X})}{\lambda_{\mu_\alpha}}\right) \quad (5)$$

For example, a distribution with $\mu_{0.5} = 128$ would be equivalent by this metric to an 8-bit random variable, denoted as $\tilde{\mu}_{0.5} = 8$. We may use units of *dits* (also called hartleys or bans) by taking base-10 logarithms instead of base-2. This represents the number of random decimal digits providing equivalent security.

3 Human choice of other 4-digit sequences

To the best of the authors’ knowledge, no dataset of real banking PINs has ever been made public. However, public datasets have recently become available for two other sources of human-chosen secret 4-digit sequences.

RockYou The leak of 32 million textual passwords from the social gaming website RockYou in 2009 has proved invaluable for password research [21]. We extracted all consecutive sequences of exactly 4 digits from the RockYou passwords. There were 1,778,095 such sequences; all possible 4-digit sequences occurred. 1234 was the most common with 66,193 occurrences (3.7%), while 8439 was the least common with 10 occurrences (0.0006%).

Though these sequences occurred as part of longer strings, a manual inspection of 100 random passwords which include a 4-digit sequence identified only 3 with an obvious connection between the digits and the text (`feb1687`, `classof2007` and `2003chevy`), suggesting that digits and text are often semantically independent. Users also show a particular affinity for 4-digit sequences, using them more significantly more often than 3-digit sequences (1,599,959) or 5-digit sequences (497,791).

iPhone Our second dataset was published (in aggregate form) in June 2011 by Daniel Amitay, an iPhone developer who deployed a screen locking mechanism which requires entering a 4-digit sequence to unlock. This dataset was much smaller, with 204,508 PINs. It doesn’t support reliable estimates of low-frequency




Fig. 1. The distribution of 4-digit sequences within RockYou passwords. Each cell shows the frequency of an individual sequence, a proxy for PIN popularity.

PINs, as 46 possible PINs weren't observed at all. 1234 was again the most common, representing 4.3% of all PINs. The screen unlock codes were entered using a square number pad very similar to standard PIN-entry pads. Geometric patterns, such as PINs consisting of digits which are adjacent on the keypad, were far more common than in the RockYou sequences.

Plotting the RockYou distribution in a 2-dimensional grid (Figure 1) highlights some basic factors influencing popularity. The most prominent features are the stripe of recent years and the range of calendar dates in MMDD and DDMM format, which trace the variation in lengths of each month. Many other features, such as a diagonal line of PINs with the same first and last two digits, and a horizontal line of PINs ending in 69, can be clearly seen.

To quantitatively measure important factors in PIN selection, we performed linear regression on each distribution with a number of human-relevant functions of each PIN as regressors. The datasets were well suited to this analysis, with nearly 10,000 samples of the response variable (the frequencies of each PIN). The assumption of a linear model simply means that the population can be divided into distinct groups of users employing different PIN selection strategies, such as choosing specific date formats or geometric patterns.

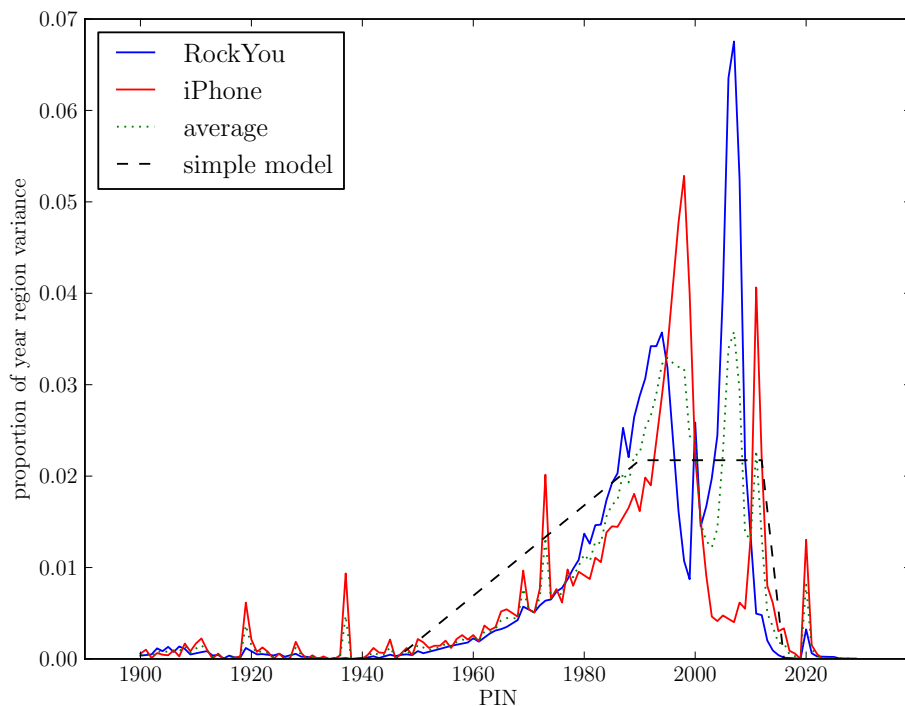


Fig. 2. Probability of 4-digit years from 1900–2025 in the RockYou dataset. Some outliers demonstrate confounding factors: 1937 and 1973 represent the four-corners of a numeric keypad, 1919 and 2020 are repeated pairs of digits, and 1969 demonstrates users’ affinity for the number 69.

Our process for identifying relevant input functions was iterative: we began with none, producing a model in which each PIN is equally likely, and progressively added functions which could explain the PINs which were the most poorly fit. We stopped at the point when we could no longer identify intuitive functions which increased the fit of the model as measured by the adjusted coefficient of determination \bar{R}^2 , which avoids bias in favour of extra input functions.

We were cautious to avoid over-fitting the training datasets, particularly for PINs which represent recent years, shown in Figure 2. The popularity of recent years has peaks between the current year and the year 1990, this range probably represents recent events like graduations or marriages (or perhaps registration). There is steady decline for older years, likely due to the drop-off in frequency of birthdays and events which are still memorable. Due to the large fluctuations for recent years in both datasets, and a possibly younger demographic for both datasets compared to the general population, we used a biased model for the popularity of different years in PIN selection: constant popularity for all years in the past 20 years, and linear drop-offs for years from 20–65 years in the past, and for 5 years into the future. This model, plotted in Figure 2, was used for PINs representing 4-digit years directly as well as DMY and MMY PINs.

factor	example	RockYou	iPhone	surveyed
date				
DDMM	2311	5.26	1.38	3.07
DMYY	3876	9.26	6.46	5.54
MMDD	1123	10.00	9.35	3.66
MMYY	0683	0.67	0.20	0.94
YYYY	1984	33.39	7.12	4.95
<i>total</i>		58.57	24.51	22.76
keypad				
adjacent	6351	1.52	4.99	—
box	1425	0.01	0.58	—
corners	9713	0.19	1.06	—
cross	8246	0.17	0.88	—
diagonal swipe	1590	0.10	1.36	—
horizontal swipe	5987	0.34	1.42	—
spelled word	5683	0.70	8.39	—
vertical swipe	8520	0.06	4.28	—
<i>total</i>		3.09	22.97	8.96
numeric				
ending in 69	6869	0.35	0.57	—
digits 0-3 only	2000	3.49	2.72	—
digits 0-6 only	5155	4.66	5.96	—
repeated pair	2525	2.31	4.11	—
repeated quad	6666	0.40	6.67	—
sequential down	3210	0.13	0.29	—
sequential up	4567	3.83	4.52	—
<i>total</i>		15.16	24.85	4.60
random selection	3271	23.17	27.67	63.68

Table 2. Results of linear regression. The percentage of the variance explained by each input function is shown for the RockYou and iPhone datasets. The final column shows estimates for the prevalence of each category from our user survey.

After fixing the year model, we removed the range of years from the regression model to avoid skewing the model’s estimation of other parameters to correct for the intentionally weakened model of the year distribution. We similarly added single-element input functions for 1234, 0000, 1111, and 2580 to avoid omitted-variable bias caused by these significant outliers.

The complete results of our final model with 25 input functions are shown in Table 2. All of the input functions were binary, except for years, calendar dates (in which Feb. 29th was discounted), and words spelled on a keypad.⁷ All of the input functions we chose contributed positively to the probability of a PIN being selected, making it plausible to interpret the weight assigned to each input function as the proportion of the population choosing a PIN by each method. The intercept term fits this interpretation naturally as the proportion of users

⁷ We used the distribution of four-letter passwords in the RockYou dataset to approximate words used in spelled-out PINs. ‘love’ was the most common 4-letter password by a large margin, and its corresponding PIN 5683 was a significant outlier.

choosing a random PIN. This simple model was able to fit both distributions quite accurately: the coefficient of determination \bar{R}^2 was 0.79 for the RockYou dataset and 0.93 for the iPhone dataset. Under the conventional interpretation, this means the model explained 79% and 93% of the variation in PIN selection.

Support for our model also comes from its accurate approximation of the source data’s guessing statistics seen in Table 1. The model consistently provides an over-approximation by about 0.2–0.3 bits (< 0.1 dit) indicating that the inaccuracy is mainly due to missing some additional sources of skew in the PIN distribution. This is acceptable for our purposes, as it will enable us to estimate an upper bound on the guessing difficulty of the banking PINs.

4 Surveying banking PIN choices

The low frequency of many PINs in the RockYou dataset means a survey of hundreds of thousands of users would be needed to observe all PINs. Additionally, ensuring that users feel comfortable disclosing their PIN in a research survey is difficult. We addressed both problems by asking users only if their PINs fall into the generic classes captured by our regression model.

We deployed our survey online using the Amazon Mechanical Turk platform, a crowd-sourcing marketplace for short tasks. The study was advertised to US-based ‘workers’ as a “Short research survey about banking security” intended to take five minutes. We deliberately displayed the University of Cambridge as the responsible body to create a trust effect. To reduce the risk of re-identification, no demographic or contact information was collected. The design was approved by the responsible ethics committee at the University of Cambridge.

The survey was piloted on 20 respondents and then administered to 1,351 respondents. 1,337 responses were kept after discarding inconsistent ones.⁸ Respondents were rewarded between US \$0.10–0.44 including bonuses for complete submission and thoughtful feedback. Repeated participation was prohibited.

4.1 PIN usage characteristics

The 1,177 respondents with a numeric banking PIN were asked a series of questions about their PIN usage. A summary of the question phrasing and responses is provided in Appendix A. A surprising number (about 19%) of users rarely or never use their PIN, relying on cash or cheques and in-person interaction with bank tellers. Several participants reported in feedback that they distrust ATM security to the point that they don’t even know their own PINs. Many others stated that they prefer signature verification to typing in their PIN. However, 41% of participants indicated that PINs were their primary authentication method for in-store payments, with another 16% using PINs or signatures equally often. Of these users, nearly all (93%) used their PINs on at least a weekly basis.

Over half of users (53%) reported sharing their PIN with another person, though this was almost exclusively a spouse, partner, or family member. This

⁸ It is common practice on Mechanical Turk tasks to include carefully-worded “test questions” to eliminate respondents who have not diligently read the instructions.

is consistent with a 2007 study which found that about half of online banking users share their passwords with a family member [18]. Of the 40% of users with more than one payment card, over a third (34%) reported using the same PIN for all cards. This rate is lower than that for online passwords, where the average password is reused across six different sites [11]. The rate of forgotten PINs was high, at 16%, although this is again broadly consistent with estimates for online passwords, where about 5% of users forget their passwords every 3 months at large websites [11]. Finally, over a third (34%) of users re-purpose their banking PIN in another authentication system. Of these, the most common were voicemail codes (21%) and Internet passwords (15%).

4.2 PIN selection strategies

We invited the 1,108 respondents with a PIN of exactly 4 digits to identify their PIN selection method. This was the most sensitive part of the survey, and users were able to not provide this information without penalty, removing a further 27% of respondents and leaving us with 805 responses from which to estimate PIN strength. We presented users with detailed descriptions and examples for each of the selection strategies identified in our regression model. Users were also able to provide free-form feedback on how they chose their PIN. The aggregated results of our survey are shown alongside our regression model in Table 2.

The largest difference between our survey results and the regression models was a huge increase in the number of random and pseudo-random PINs: almost 64% of respondents in our survey, compared to 23% and 27% estimated for our example data sets. Of these users, 63% reported that they either used the PIN initially assigned by their bank or a PIN assigned by a previous bank.⁹ Another 21% reported the use of random digits from another number assigned to them, usually either a phone number or an ID number from the government, an employer, or a university (about 30% for each source).¹⁰

Of users with non-random PINs, dates were by far the largest category, representing about 23% of users (comparable to the iPhone data and about half the rate of the RockYou data). The choice of date formats was similar to the other datasets with the exception of 4-digit years, which were less common in our survey. We also asked users about the significance of the dates in their PINs: 29% used their own birth date, 26% the birth date of a partner or family member, and 25% an important life event like an anniversary or graduation.

Finally, about 9% of users chose a pattern on the keypad, and 5% a numeric pattern such as repeated or sequential digits. Our sample size was insufficient to provide an accurate breakdown of users within these categories.

⁹ We explored the possibility that some of our participants kept their initial PIN simply because they rarely or never used their card, but the rate was statistically indistinguishable for users using their PIN at least once per week.

¹⁰ While reusing identification numbers and phone numbers in PINs may open a user to targeted attacks, they should appear random to a guessing attacker.

guessing scenario	H_1	\tilde{G}	$\tilde{\mu}_{0.5}$	λ_3	λ_6
baseline	12.90	12.83	12.56	1.44%	1.94%
with blacklist	13.13	12.95	12.79	0.12%	0.24%
known birth date	12.57	12.80	12.49	5.52%	8.23%
blacklist, known birth date	12.85	12.92	12.75	5.11%	5.63%
random 4-digit PIN	13.29	13.29	13.29	0.03%	0.06%

Table 3. Guessing metrics for banking PINs, using the model computed from our survey and regression results on the iPhone dataset.

5 Approximating banking PIN strength

Using our survey data and regression model we estimated the distribution of banking PINs for our survey population. This was straightforward for random PINs and PINs based on dates. Within the other two categories we used the sub-distribution from the iPhone dataset due to lack of sufficient sample size.

Statistics for our best estimation are shown in Table 3. By any of the aggregate metrics $\tilde{\mu}_{0.5}$, \tilde{G} , or H_1 , the strength is actually quite good—between 12.6 and 12.9 bits (3.8–3.9 dits), close to the maximum possible. In other words, if an attacker can try many PINs for a targeted card, the introduction of human choice does not significantly reduce security compared to randomly-assigned PINs.

Banking PINs appear considerably more vulnerable against marginal guessing attacks. As noted in Table 3, an attacker with 3 guesses will have a $\lambda_3 = 1.4\%$ chance of success and an attacker with 6 guesses a $\lambda_6 = 1.9\%$ chance of success, equivalent to $\tilde{\lambda}_6 = 8.3$ bits of security (2.5 dits). This is significantly better than the estimates based on the RockYou or iPhone distributions (Table 1), for which $\lambda_6 > 10\%$. The optimal guessing order is 1234 followed by 1990–1986.

5.1 Known birth date guessing

Given the large number of users who base their PIN on their birth date (nearly 7% in total, or 29% of those using some type of date), we evaluated the success of an attacker who can leverage a known birth date, for example if a card is stolen in a wallet along with an identification card. The exact effects vary slightly with the actual birth date: if variants of the date also correspond to common PINs such as 1212, the attacker’s success rate will be higher. We calculated guessing probabilities for all dates from 1960–1990 and report results for the median date of June 3, 1983. In this scenario, the attacker’s optimal strategy shifts to guessing, in order, 1983, 6383, 0306, 0603, 1234, and 0683. As seen in Table 3, the attacker benefits considerably from this knowledge: λ_6 increases to 8.2%, providing only $\tilde{\lambda}_6 = 6.2$ bits (1.9 dits) of security.

5.2 Effectiveness of blacklisting

Assuming that users with a blacklisted PIN will be re-distributed randomly according to the rest of the distribution (as assumed in [21]), the effects of blacklisting the top 100 PINs are substantial— λ_6 drops to 0.2%.¹¹ This is equivalent

¹¹ The optimal blacklist suggested by our model is given in Appendix B.

guessing scenario	number of stolen cards				exp.
	1	2	3	4	
baseline	1.9%	2.9%	3.9%	4.9%	2.5%
with blacklist	0.2%	0.5%	0.7%	0.9%	0.4%
known birth date	8.2%	9.7%	10.3%	10.9%	8.9%
blacklist, known birth date	5.6%	6.0%	6.2%	6.4%	5.8%
random 4-digit PIN	0.1%	0.1%	0.2%	0.2%	0.1%

Table 4. Probability of a successful attack given multiple cards from one user. The final column is an expected value given the observed rate of card ownership.

to $\tilde{\lambda}_6 = 11.6$ bits (3.9 dits) of security, indicating that a very small blacklist may eliminate most insecurity due to human choice. Unfortunately, as seen in Table 3 and Table 4, blacklisting is much less effective against known birth date attacks, only reducing λ_6 to 5.1% ($\tilde{\lambda}_6 = 6.9$ bits/2.1 dits). With a reasonable blacklist, it is only possible to block the YYYY format, leaving an attacker to try DDMM, MMDD, and so on; preventing this would require user-specific blacklists.

5.3 Expected value of a stolen wallet

We calculate the guessing probability of a thief with multiple stolen cards, for example from an entire wallet or purse, in Table 4. Though most of our surveyed users own only one card with a PIN, on expectation stealing a wallet instead of a single card raises a thief’s guessing chances by over a third. Our survey results suggest that virtually all payment card users (99%) carry documentation of their birth date alongside their card.¹² Thus, we conclude that a competent thief will gain use of a payment card once every 11–18 stolen wallets, depending on the proportion of banks using a denied PIN list.

6 Concluding remarks

The widespread security role assigned to 4-digit PINs is a historical accident which has received surprisingly little scrutiny. While complete analysis is impossible without access to a huge list of real banking PINs, it appears that user choice of banking PINs is not as bad as with other secrets like passwords. User management of PINs is also comparatively good, with lower rates of reuse and sharing and many users reporting serious thought about PIN security. However, the skew introduced by user choice may make manual guessing by thieves worthwhile—a lost or stolen wallet will be vulnerable up to 8.9% of the time in the absence of denied PIN lists, with birthday-based guessing the most effective strategy. Blacklisting appears effective only if a thief doesn’t know the user’s date of birth (or users stop using this to choose their PIN). We advise users not to use PINs based on a date of birth, and those banks which do not currently employ blacklists to immediately do so. Still, preventing birthday-based guessing requires a move away from customer-chosen PINs entirely.

¹² The prevalence of carrying ID may vary by locale. In 24 US states carrying ID is legally required. In the UK, carrying ID is not required and fewer citizens carry it.

References

1. *EMV Integrated Circuit Card Standard for Payment Systems version 4.2*. EMVco, 2008.
2. Issuer PIN Security Guidelines. Technical report, VISA, November 2010.
3. *ISO 9564:2011 Financial services – Personal Identification Number (PIN) management and security*. International Organisation for Standardisation, 2011.
4. B. Batiz-Lazo and R. J. Reid. The Development of Cash-Dispensing Technology in the UK. *IEEE Annals of the History of Computing*, 33:32–45, 2011.
5. M. Bond and P. Zieliński. Decimalisation table attacks for PIN cracking. Technical Report UCAM-CL-TR-560, University of Cambridge, Jan. 2003.
6. J. Bonneau, M. Just, and G. Matthews. What’s in a name? Evaluating statistical attacks against personal knowledge questions. *FC ’10: The Fourteenth International Conference on Financial Cryptography and Data Security*, 2010.
7. S. Boztas. Entropies, Guessing, and Cryptography. Technical Report 6, Department of Mathematics, Royal Melbourne Institute of Technology, 1999.
8. W. E. Burr, D. F. Dodson, and W. T. Polk. Electronic Authentication Guideline. *NIST Special Publication 800-63*, April 2006.
9. C. Cachin. *Entropy measures and unconditional security in cryptography*. PhD thesis, ETH Zurich, 1997.
10. S. Drimer, S. J. Murdoch, and R. Anderson. Optimised to Fail: Card Readers for Online Banking. *FC ’09: The Thirteenth International Conference on Financial Cryptography and Data Security*, 2009.
11. D. Florencio and C. Herley. A large-scale study of web password habits. In *WWW ’07: Proceedings of the 16th International Conference on World Wide Web*, pages 657–666, New York, NY, USA, 2007. ACM.
12. A. Ivan and J. Goodfellow. Improvements in or relating to Customer-Operated Dispensing Systems . UK Patent #GB1197183, 1966.
13. M. Kuhn. Probability Theory for Pickpockets—ec-PIN Guessing . Technical report, Purdue University, 1997.
14. J. L. Massey. Guessing and Entropy. In *Proceedings of the 1994 IEEE International Symposium on Information Theory*, page 204, 1994.
15. R. Morris and K. Thompson. Password security: a case history. *Commun. ACM*, 22(11):594–597, 1979.
16. S. J. Murdoch, S. Drimer, R. Anderson, and M. Bond. Chip and PIN is Broken. *Security and Privacy, IEEE Symposium on*, 0:433–446, 2010.
17. J. O. Pliam. On the Incomparability of Entropy and Marginal Guesswork in Brute-Force Attacks. In *Progress in Cryptology-INDOCRYPT 2000*, 2000.
18. S. Singh, A. Cabraal, C. Demosthenous, G. Astbrink, and M. Furlong. Password Sharing: Implications for Security Design Based on Social Practice. In *CHI ’07: Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, pages 895–904, New York, NY, USA, 2007. ACM.
19. E. Spafford. Observations on Reusable Password Choices. In *Proceedings of the 3rd USENIX Security Workshop*, 1992.
20. P. C. van Oorschot and J. Thorpe. On Predictive Models and User-Drawn Graphical Passwords. *ACM Trans. Inf. Syst. Secur.*, 10(4):1–33, 2008.
21. 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, CCS ’10*, pages 162–175, New York, NY, USA, 2010. ACM.

A Survey presentation and results

The following is a summary of questions about user PIN management. The complete survey, including questions about PIN selection, is available online at: http://preibusch.de/publications/pin_survey/.

Do you regularly use a PIN number with your payment cards? ($N = 1337$)

yes, a 4-digit PIN	yes, a PIN of 5+ digits	no
1108 (82.9%)	69 (5.2%)	160 (12.0%)

When making purchases in a shop, how do you typically pay? ($N = 1177$)

I use my payment card and key in my PIN	477 (40.5%)
I use my payment card and sign a receipt	357 (30.3%)
I use my payment card with my PIN or my signature equally often	184 (15.6%)
I normally use cash or cheque payments and rarely use payment cards	159 (13.5%)

Overall, how often do you type your PIN when making a purchase in a shop? And how often do you type your PIN at an ATM/cash machine? ($N = 1177$)

	shop		ATM	
Multiple times per day	81 (6.9%)	14 (1.2%)	14 (1.2%)	
About once per day	117 (9.9%)	19 (1.6%)	19 (1.6%)	
Several times a week	342 (29.1%)	118 (10.0%)	118 (10.0%)	
About once per week	241 (20.5%)	384 (32.6%)	384 (32.6%)	
About once per month	113 (9.6%)	418 (35.5%)	418 (35.5%)	
Rarely or never	283 (24.0%)	224 (19.0%)	224 (19.0%)	

How many payment cards with a PIN do you use? ($N = 1177$)

1	2	3	4	5	6
708 (60.2%)	344 (29.2%)	89 (7.6%)	23 (2.0%)	11 (0.9%)	2 (0.2%)

Median: 1, Mean: 1.5

If you have more than one payment card which requires a PIN, do you use the same PIN for several cards? ($N = 469$)

yes	no
161 (34.3%)	308 (65.7%)

Have you ever changed the PIN associated with a payment card? ($N = 1177$)

Never	Yes, when I initially received the card	Yes, I change periodically
591 (50.2%)	376 (31.9%)	210 (17.8%)

Have you ever forgotten your PIN and had to have your financial institution remind you or reset your card?(N = 1177)

yes	no
186 (15.8%)	991 (84.2%)

Have you ever shared your PIN with another person so that they could borrow your payment card?(N = 1177)

spouse or significant other	475	(40.4%)
child, parent, sibling, or other family member	204	(17.3%)
friend or acquaintance	40	(3.4%)
secretary or personal assistant	1	(0.1%)
any	621	(52.8%)

Have you ever used a PIN from a payment card for something other than making a payment or retrieving money?(N = 1177)

password for an Internet account	180	(15.3%)
password for my computer	94	(8.0%)
code for my voicemail	242	(20.6%)
to unlock the screen for mobile phone	104	(8.8%)
to unlock my SIM card	29	(2.5%)
entry code for a building	74	(6.3%)
any	399	(33.9%)

Do you carry any of the following in your wallet or purse?(N = 415)¹³

driver's license	377	(90.8%)
passport or government ID card	68	(16.4%)
social security or other insurance card	155	(37.3%)
school or employer ID listing date of birth	23	(5.5%)
other document listing date of birth	78	(18.7%)
any item with date of birth	411	(99.0%)

B Suggested blacklist

According to our computed model, the following blacklist of 100 PINs is optimal: 0000, 0101–0103, 0110, 0111, 0123, 0202, 0303, 0404, 0505, 0606, 0707, 0808, 0909, 1010, 1101–1103, 1110–1112, 1123, 1201–1203, 1210–1212, 1234, 1956–2015, 2222, 2229, 2580, 3333, 4444, 5252, 5683, 6666, 7465, 7667.

C Acknowledgements

We thank Mike Bond, Andrew Lewis, Saar Drimer, Steven Murdoch, and Richard Clayton for helpful discussions about PIN security, Bernardo Batiz-Lazo for comments about ATM history, Alastair Beresford for assistance with survey design, and Daniel Amitay for sharing data. We also thank Andra Adams, Jon Anderson, Alexandra Bowe, Omar Choudary, William Swan Helvestine, Markus Kuhn, Niraj Lal, Will Moreland, Frank Stajano, Paul van Oorschot and Robert Watson for help identifying banking practices around the world. Joseph Bonneau was supported by the Gates Cambridge Trust during this research.

¹³ This question was sent to a random subset of respondents after the main survey.