I/S: A JOURNAL OF LAW AND POLICY FOR THE INFORMATION SOCIETY

# Monetary Forgery in the Digital Age: Will Physical-Digital Cash Be a Solution?[1]

NICOLAS CHRISTIN[2], ALESSANDRO ACQUISTI[3], BRYAN PARNO[4] AND ADRIAN PERRIG[5]

*Abstract: Despite cryptographic breakthroughs in the area of digital cash and the rapid advance of information technology, physical cash remains the dominant currency: it is easy to use and its exchanges are largely independent of computing devices. However, physical cash is vulnerable to rising threats such as high quality, government-mandated forgeries. Can a hybrid of physical and digital cash protect more effectively against these threats? We discuss the rise of high-quality counterfeits and review technological solutions to thwart such threats. Specifically, we study mechanisms to combine physical cash with digital cash to remove their respective shortcomings and obtain their combined advantages. The mechanisms range from cryptographic signatures embedded in 2-D barcodes to online verification systems assisted by physical one-way functions. Notably, we compare these different proposals by looking at them through the prism of economics, and examining their cost and benefit trade-offs.*

---

[1] A preliminary, abbreviated version appeared in the Proceedings of the Twelfth International Conference on Financial Cryptography and Data Security (FC'08). Cozumel, Mexico. January 28–31, 2008.

[2] Carnegie Mellon University, Information Networking Institute

[3] Carnegie Mellon University, Heinz School of Public Policy and Management

[4] Carnegie Mellon University, Department of Electrical and Computer Engineering & Microsoft Research

[5] Carnegie Mellon University, Department of Electrical and Computer Engineering & Carnegie Mellon University, Department of Engineering and Public Policy

## I. INTRODUCTION

Counterfeiting money is arguably as old as minting money. Fake coins have been discovered dating back to the 4th century B.C.[6] During World War II, both Allied and Axis powers contemplated flooding their adversaries with forged banknotes.[7] While the Allies eventually decided against implementing the attack,[8] Nazi Germany went ahead with the plan, and produced high quality forgeries of British Pound banknotes. While its effects were not necessarily negligible, the attack was hardly a success in terms of monetary destabilization.

More recently, circumstantial evidence of a nation-state issuing large amounts of nearly perfect counterfeit US dollars ("supernotes") may have surfaced.[9] But where Nazi Germany's counterfeiting strategies during World War II relied mostly on master counterfeiters and their human labor (highly skilled, but time consuming), today's digital printing technologies make it possible to produce vast amounts of counterfeited currencies in limited periods of time. From a security engineering standpoint, the mere possibility of modern, high-quality forgeries produced by hostile governments or sophisticated third-party attackers with access to considerable resources calls for a reevaluation of the traditional threat model used to design anti-counterfeiting techniques, both in terms of technology available, and potential magnitude of the losses. Regardless of whether the counterfeiters are actually employed by a government, or transnational criminal syndicate, the appearance of what we term *high-quality monetary forgery* marks a significant departure from

---

[6] G. Giovanelli et al.,"A Puzzling Mule Coin from the Parabita Hoard: a Material Characterization" (paper presented at the Cavallino Archaeometry Workshop, Lecce, Italy, May 2006).

[7] L. Malkin. Krueger's Men: The Secret Nazi Counterfeit Plot and the Prisoners of Block 19. Little, Brown and Company, 2006.

[8] Both on political and economic grounds. A May 1942 exchange between Sigismund David Waley (at the British Treasury) and John Maynard Keynes (the famous economist) can be found in documents made available by [26]. In the exchange, the two gentlemen reject the plan, "as the amount which could be introduced would not make any appreciable difference." *See* http://www.lawrencemalkin.com/kruegers-men-the-secret-documents-2-3-1.html.

[9] S. Mihm, "No Ordinary Counterfeit," *New York Times Magazine*, July 23, 2006, 36.

traditional forgery (e.g., that perpetrated by smaller organized crime outfits) in scale, motivation, and perception.

We define high-quality as follows: we assume that the counterfeiting entity has access to manufacturing resources and capabilities that can be considered equivalent in quality and production levels to that of the national bank whose currency is being faked. As a result, the counterfeited notes are indistinguishable from the legitimate currencies, unless sophisticated (and therefore very costly) forensic analysis is employed.[10] While these counterfeits may simply be used to increase the purchasing power of the entity producing the forgeries, the forged bills may also be used to finance hostile activities, such as weapons purchases, bribing of spies, or terrorism sponsorship. Consequently, targeted countries may be willing to consider relatively expensive defenses against such high-quality forgeries.

On the other hand, the development of the Internet has presided over a considerable expansion of infrastructures for online transactions and has fostered large-scale deployment of security technologies such as cryptographic verification. By making feasible a host of online verification schemes, increased connectivity may be an important asset in the design of novel anti-counterfeit defenses.

We offer the following contributions with this paper: we investigate the economic impact of counterfeiting attacks and we examine the feasibility of a set of technological solutions, both offline and online, against counterfeiting, given the economic constraints imposed by the volume of banknotes in production and circulation. In particular, we discuss the main technical and economic challenges related to the design and deployment of possible countermeasures against high-quality monetary forgery, and debate whether technology-based solutions, such as electronic cash, may be an answer to those challenges.

One of the main motivating factors for this study is that, despite major developments in paperless currency over the past decade, physical cash remains widely used throughout the world. As such, the entry barrier for adoption of alternate proposals is extremely high. An appealing aspect of physical cash is that people can trade it without the assistance of computing devices. People expect that simple visual and tactile inspection reveals fake bills. Physical cash can survive extreme situations: it can endure a cycle in a washing machine and it

---

[10] S. Mihm, "No Ordinary Counterfeit," *New York Times Magazine*, July 23, 2006, 36.

can survive extreme temperatures that would render any smartcard unusable. Although not perfectly anonymous[11], physical cash, especially smaller and widely circulated bills, provides a reasonable level of privacy.

On the other hand, paperless, cryptographic digital cash offers numerous benefits, and provides two key advantages over physical money. First, an adversary cannot forge digital cash, assuming the security of the cryptographic mechanisms and the secrecy of the associated cryptographic information are preserved. Second, replication of digital cash is easy, so that one can easily safeguard against loss or theft of digital cash through digital backups.

The idea to use digital cash as a solution against money counterfeiting is not new, and was in fact suggested as soon as the first practical digital cash schemes were available.[12] This has received considerable interest from the technical community.[13] Counterfeit resilience also spurred a large body of research as one application of quantum cryptography,[14] although current quantum cryptography implementations are still far from being practical. More recent efforts (e.g., R. Balan et al., "mFerio: The design and evaluation of a peer to-peer mobile payment system"[15]) argue that with the ubiquity of cellular phones and PDAs, mobile devices can greatly facilitate adoption of digital cash by the masses.

However, the transaction, coordination, and social costs associated with any large scale switch to digital cash explain why, in spite of the advance of cell phones and credit cards, we are still far from becoming a cashless society, especially in many developing nations.[16] It may be more beneficial for an economy to preserve the

---

[11] D. Kügler, "On the Anonymity of Banknotes" (paper presented at the 4th International Workshop on Privacy Enhancing Technologies (PET'04), Toronto, Canada, May 2004).

[12] L. Malkin, *Krueger's Men: The Secret Nazi Counterfeit Plot and the Prisoners of Block 19*, Little, Brown and Company, 2006.

[13] *See* Patiwat Panurach, "Money in electronic commerce: digital cash, electronic fund transfer, and ecash." *Commun. ACM*, 39(6):45–50, 1996.

[14] *See* S. Wiesner, "Conjugate Coding," *SIGACT News*, 15(1):78–88, 1983.

[15] R. Balan et al., "mFerio: The design and evaluation of a peer to-peer mobile payment system" (paper presented at the Seventh ACM/USENIX Annual International Conference on Mobile Systems, Applications and Services (MobiSys '09), Krakow, Poland, June 2009).

[16] Even developed nations at the forefront of technological advances rely heavily on physical cash. For instance, for a variety of reasons discussed by Mann [27], Japanese society has traditionally been relatively reluctant to use credit and debit card transactions.

appealing aspects of physical cash, including its ubiquitous deployment, and combine those with the advantages of digital cash—in essence, a kind of physical-digital money.

Indeed, a few hybrid solutions coupling physical security with cryptographic verification have been suggested.[17] Each has its own specific trade-offs. By embedding an easily verifiable cryptographic value in regular bills,[18] the issuing government can combine physical and digital cash without requiring drastic changes to the underlying existing monetary infrastructure. However, devising such bills, or physical-digital cash, also leads to a number of design trade-offs between the security properties achieved, the technological complexity involved, and the economic costs incurred. In this paper, we explore the trade-offs of these and alternative solutions in search of deployable anti-counterfeiting techniques, noting that recent advances in large-scale communications and databases can provide additional layers of defense against forgery.

After surveying existing anti-counterfeiting technological solutions and related work on digital cash in Part II, we use an economic framework in Part III to inform design requirements for counterfeit-resistant bills, which is examined in Part IV. We contrast the advantages and disadvantages of several schemes in Part V. These schemes offer various levels of protection against both basic theft and attempts at high-quality forgery. We then analyze general security threats against physical-digital cash in Part VI, and offer some final remarks in Part VII.

---

[17] A. Acquisti et al., "Countermeasures against government-scale monetary forgery." In *Proceedings of the Twelfth International Conference on Financial Cryptography and Data Security (FC'08)*, pages 262–266, Cozumel, Mexico, January 2008; [22] H. Hoshino, et al. "Object to be checked for authenticity and a method for manufacturing the same," February 1997. US Patent nr. 5,601,931; [37] Ravikanth Pappu et al., "Physical one-way functions." *Science*, 297(5589):2026–2030, 2002.; [45] G. J. Simmons. " Identification of data, devices, documents and individuals." In *Proc. 25th Ann. Intern. Carnahan Conference on Security Technology*, pages 197–218, Taipei, Taiwan, ROC, October 1991. IEEE.

[18] In this paper we focus our discussion on bills, although the principles we present could be translated to coins as well. Given the lower economic value of coins and their high cost of production, counterfeiting coins is usually not effective.

## II. RELATED WORK

Proposals to combat monetary forgery by building protection into the currency can be classified into three groups: physical protection of currency, digital alternatives to cash, and hybrid approaches, which attempt to combine physical features with digital enhancements.

### A. PHYSICAL PROTECTION

In the area of physical protection against counterfeits, each currency-printing nation has developed its own secret techniques. However, a number of public features enable people to visually inspect and verify the authenticity of each bill. For example, the U.S. Bureau of Printing and Engraving publishes details about some of the features of new U.S. dollars, such as color-shifting ink, a new watermark, a metallic security thread, and the use of micro print.[19] Euro bank notes also provide numerous security features, including raised print, watermarks, a security thread, see-through numbers, holograms, a glossy stripe, a color-changing number, and UV-visible features.[20] The most valuable bank note in the world, the 1000 Swiss Franc bill, includes a kinegram, an irodin number, a watermark, UV-visible features, numbers visible only under oblique incident light (the Kipp effect), and the use of copper print, micro perforation, and optically variable ink.[21] Recent research showed that even for common types of paper, random, natural imperfections occurring in the paper texture make it possible to authenticate documents,[22] which could, of course, be useful for counterfeit resilience.

---

[19] U.S. Bureau of Engraving and Printing. Accessed August 30, 2010. http://www.bep.treas.gov/document.cfm/18/106.

[20] Euro Banknotes – Security Features" European Central Bank, accessed August 30, 2010, http://www.ecb.int/bc/banknotes/ security/html/index.en.html.

[21] Schweizerische Nationalbank, Banque Nationale Suisse. Die aktuelle banknotenserie. http://www.snb. ch/d/banknoten/aktuelle_serie/aktuelle_serie.html.

[22] William Clarkson et al., "Fingerprinting blank paper using commodity scanners." (paper presented at the *IEEE Symposium on Security and Privacy*, May 2009.)

## B. DIGITAL ALTERNATIVES TO CASH

Many researchers have proposed and studied implementations of digital cash schemes. Asokan et al. [23] provide an overview article of electronic payment systems. The key idea is to provide an alternative to existing cash-based payment systems that dispense paper money to users, and replace it with cryptographically-verifiable electronic tokens.

Based on seminal work on blind signatures,[24] one line of research focuses on cryptographic digital cash systems.[25] Similarly, several micropayment systems have also been proposed to pay for very small amounts.[26] The core motivation for the line of work in micropayments is not to increase transaction security, but instead to provide economically efficient alternatives to credit card payments (which incur relatively high processing fees) for small transactions. In terms of security primitives, rather than trying to defeat existing counterfeiting operations, micropayments focus on providing security guarantees to emulate existing currency usage—for instance, by making it impossible for a specific user to spend the same coin twice on different transactions.

Another line of research focuses on trusted hardware-based payment systems [27] such as electronic wallets. In particular, the

---

[23] N. Asokan et al. "State of the Art in Electronic Payment Systems" *Advances in Computers* 43 (2000): 425-449.

[24] David Chaum, "Blind Signatures for Untraceable Payments" (paper presented at CRYPTO'82, 1982).

[25] *See* D. Chaum, "Untraceable electronic cash." In *Proc. CRYPTO'88*, pages 319–327, 1988.

[26] Benjamin Cox, et al., "NetBill Security and Transaction Protocol" (paper presented at the 1st USENIX Workshop on E-Commerce, New York, NY, 1995); Steve Glassman et al., "The MilliCent Protocol for Inexpensive Electronic Commerce" (paper presented at.WWW'95, Boston, MA, December 1995); Ronald Rivest, "Electronic Lottery Tickets as Micropayments," Presented at the International Conference on Financial Cryptography and Data Security 1997, Anguilla, BWI, February 1997, pages 307–314.; Ronald Rivest and Adi Shamir, "PayWord and MicroMint: Two Simple Micropayment Schemes," Presented at the Int'l Workshop on Security Protocols, Cambridge, UK, April 1997, pages 69 – 88.

[27] Jean-Paul Boly et al.,"The ESPRIT Project CAFE -High Security Digital Payment Systems" (paper presented at ESORICS'94, 1994); Mondex. http://www.mondex.com/mondex/home.htm; Sony Corporation. Overview of FeliCa. http://www.sony.net/Products/felica/abt/dvs. html.

Mondex system[28] and the CAFE project[29] proposed portable trusted hardware devices to store an account balance and perform electronic payments. The FeliCa card[30] proposed by Sony is a type of RFID-based electronic wallet becoming increasingly popular in East Asia. Japanese railways (JR East and West) use FeliCa for train passes, and the Octopus Card in Hong Kong uses FeliCa as a debit card. FeliCa's popularity is partly due to its seamless integration in some cell phones handsets. More recent proposals investigate how cell phones can become a digital wallet and replace cash.[31]

### C. HYBRID APPROACHES (PHYSICAL-DIGITAL CASH)

In an attempt to enhance existing bank notes, hybrid approaches include cryptographic materials within "usual" bills. The idea is to maintain compatibility with the existing monetary infrastructure, while providing electronic enhancements that can be used by parties who wish to do so to verify the genuineness of a given bill.

For instance, in 2001, the European Central Bank considered embedding RFID tags in each Euro note.[32] These tags would give each bill a unique identifier and embed cryptographic material attesting to the validity of the note. A few other proposals have attempted to couple physical security, using physical one-way functions,[33] with cryptographic verification of the bill.[34]

---

[28] Mondex, "What is Mondex". http://www.mondex.com/mondex/home.htm.

[29] Jean-Paul Boly et al., "The ESPRIT Project CAFE -High Security Digital Payment Systems" (paper presented at *ESORICS'94, 1994).*

[30] Sony Corporation. Overview of FeliCa. http://www.sony.net/Products/felica/abt/dvs.html.

[31] R. Balan et al., "mFerio: The design and evaluation of a peer-to-peer mobile payment system." (paper presented at the Seventh ACM/USENIX Annual International Conference on Mobile Systems, Applications and Services (MobiSys '09), Krakow, Poland, June 2009.

[32] Junko Yoshida, "Euro bank notes to embed RFID chips by 2005." *EE Times* (2001). http://www.eetimes.com/ story/OEG20011219S0016, December 2001.

[33] Ravikanth Pappu et al., "Physical One-way Functions," *Science*, 297(5589) (2002): 2026–2030, 2002.

[34] H. Hoshino, et al., "Object to be Checked for Authenticity and a Method for Manufacturing the ame," February 1997. US Patent No. 5,601,931; [45] G. J. Simmons. "Identification of data, devices, documents and individuals." (paper presented at the *25th Ann. Intern. Carnahan Conference on Security Technology*, pages 197–218, Taipei, Taiwan, ROC, October 1991.)

In Part V. we discuss and contrast the trade-offs associated with the above approaches, as well as novel approaches introduced in this paper. Before doing so, we consider the economic impact of forgeries, in order to better understand the technological and economic trade-offs a government faces when deploying counterfeit-resistant bills.

## III. ECONOMICS OF COUNTERFEITING

This Part highlights the economic implications of counterfeiting, focusing on large-scale, high-tech counterfeiting that new technologies have made possible. We first consider the perspective of the defender, analyzing the costs and benefits of engaging in various strategies to fight forgeries; we then consider the benefits for the attackers, and relate them to the destabilizing effects that forgeries can have on an economy. The discussion of these effects leads us to consider constraints on the costs of new countermeasures.

### A. COMBATING COUNTERFEITING

Efforts to combat counterfeiting can focus on three broad areas: policing (actively pursuing counterfeiters and their distribution channels, and reducing the incentives to engage in counterfeiting by punishing violators), building protection into the currency, and detecting counterfeits in circulation.

In the United States, anti-counterfeiting has been historically tackled by the United States Secret Service (USSS), which now operates under the Department of Homeland Security (DHS). Even though the USSS is perhaps most famous for being in charge of protection of high-ranking diplomats and heads of state, including the US President, combating financial crimes (including currency counterfeiting) is its main activity.

Hence, investigating the DHS budget can give a sense of the amounts invested by the US government into policing and detecting counterfeits in circulation. The (USSS) operates on a US $1.4 billion budget. Out of this budget, somewhere between $270 million and $300 million per year are used to fight financial crimes;[35] most of this budget is directed toward domestic and international field operations.[36]

---

[35] Department of Homeland Security. FY 2009 Budget Details, 2009. http://www.dhs.gov/xlibrary/assets/budget_fy2009.pdf.

[36] Department of Homeland Security. 2009. FY 2009 Budget Details, 2009. http://www.dhs.gov/xlibrary/assets/budget_fy2009.pdf.

It is hard to infer exactly how much of this budget is devoted to anti-counterfeiting, as the USSS also deals with online crime, money laundering, and other forms of financial crimes. However, the USSS claims they prevent between $1 billion and $4 billion per year in financial losses due to counterfeiting and other monetary crimes.[37] The USSS also reports that the amount of counterfeits in circulation represents less than 0.01% of the total currency in circulation (the actual estimate in 2007 was 0.0079%).[38]

The Federal Reserve Board budget for 2009[39] also provides some insight into the costs of building security into the currency. Current $20 and $50 dollar bills cost almost 10 cents per bill to produce. New $100 bills cost 13 cents each. On the other hand, the costs to produce $1/$2 bills are in the 5-cent/unit range. Printing represents 96% of the currency budget. Counterfeit-deterrence research is about $4.2 million per year—a small amount compared to the resources invested in policing. The research budget has however seen an increase of 15% over the past year.

### B. INCENTIVES TO ENGAGE IN COUNTERFEITING MONEY

One key motivation for embedding more physical security features in banknotes is that counterfeiters now have access to reasonably high-quality printing equipment for a fraction of the amount it used to cost. Color printers and scanners, once prohibitively expensive, can now be purchased even under a limited budget; and while such equipment can only provide low-quality counterfeits, it is worth noting that this does not deter counterfeiters. Indeed, according to the USSS, in 2001, about 39% of the $47.5 million in seized counterfeit money that entered circulation in the United States was made using computers or scanners. In 1995, the figure was less than a half percent.[40] Such an increase suggests that counterfeiters have strong

---

[37] *See* Department of Homeland Security. FY 2009 Budget Details, 2009, p. 1858. http://www.dhs.gov/xlibrary/assets/budget_fy2009.pdf.

[38] *See* Department of Homeland Security. FY 2009 Budget Details, 2009, p. 1838. http://www.dhs.gov/xlibrary/assets/budget_fy2009.pdf.

[39] "2009 New Currency Budget," The Federal Reserve Board, accessed August 30, 2010, http://www.federalreserve.gov/generalinfo/foia/2009newcurrency.htm.

[40] C. Marshall, "Paper or Plastic? Currency Making is in Flux." *New York Times*, July 14, 2002, last accessed on January 22, 2011, http://www.nytimes.com/2002/07/14/business/business-paper-or-plastic-currency-making-is-in-flux.html.

incentives to produce low-quality counterfeits: production costs are extremely low, and remain inferior to the expected returns.

To better understand the value of counterfeit money, notice that the counterfeit "supply chain" is usually more complex than a simple producer-consumer relationship. Counterfeiters are unlikely to directly inject the counterfeit bills into the market, and instead will pass them on to (a chain of) intermediaries that will exchange the counterfeit money at a discounted price for other goods, before trying to pass them onto the next element of the supply chain at a higher value.[41] Eventually, the principal who injects the currency into the market hopes to obtain face value (i.e., cash out $100 for a counterfeit $100 bill), after having spent significantly less to obtain the counterfeit (e.g., $70 for a supernote). However, the risks of getting caught are considerably higher when injecting the currency into the market than they are when exchanging the counterfeits for goods or services (provided the other party to the transaction knows and agrees that counterfeits are being used). In particular, counterfeit currency has been known to be accepted as a legitimate means of payment in the context of the drug trade. Repasky reports that counterfeits of high quality can be used at about 33 cents on the dollar in exchange for drugs.[42] The fact that counterfeit currency facilitates the drug trade provides a strong impetus to governments to try to thwart counterfeit money.

In such a context, an extremely high-quality counterfeit, i.e., a "supernote," such as those that may be produced by rogue states or large crime syndicates by using the same printing technologies employed by the respective governments, can return about 60 to 70 cents on the dollar to the counterfeiter.[43] This price not only reflects higher production costs, but also a considerably lower risk of getting caught for the party injecting the bill into the market.

---

[41] R. Perl and D. Nanto, "North Korean Counterfeiting of U.S. Currency," 2007, Congressional Research Service report RL33324. Available online at http://opencrs.com/document/RL33324/2007-01-17/.

[42] R. Repasky, "Currency and Financial Crimes in the New Millennium." In *Protection, Security, Safeguards: Practical Approaches and Perspectives*. Los Angles: Henley-Putnam University, 2000, 197 – 211.

[43] R. Perl and D. Nanto, "North Korean Counterfeiting of U.S. Currency," 2007, Congressional Research Service report RL33324. Available online at http://opencrs.com/document/RL33324/2007-01-17/.

## C. ECONOMIC IMPACT OF FORGERIES

So far, we have mostly considered criminal motives for producing counterfeits. Could there be political motives that go far beyond mere profit, in the production of forgeries? At least in principle, a large-scale forgery produced by a hostile government could increase the monetary base enough to affect the rate of inflation in the targeted nation and decrease national and international confidence in that currency, further destabilizing its exchange rate. Is this a practical threat?

The United States has, in the past, experienced periods of monetary instability caused by large amounts of forged currencies in their money supply. For instance, it is estimated that at the end of the Civil War, between one-third and one-half of U.S. currency was counterfeit.[44] Yet, market economies can be resilient to shocks in their monetary supplies: less than ten years after the end of the war, the amount of forged U.S. currency in circulation had been drastically reduced.

According to the so-called quantity theory of money,[45] the effect on inflation of an increase in the money supply can be estimated by comparing the rate of money growth to the change in money velocity (the number of times currency turns over in a year) and the change in real gross domestic product. Assuming, for instance, that Gross Domestic Product (GDP) and velocity are not changing, an increase of $x$% in the money supply will cause an approximate increase of $x$% in the inflation rate. To give a sense of these dynamics using our previous example, between 1861 and 1864 the Confederate money supply increased 11.5 times, causing commodity prices to increase 28 times.[46]

As of February 2006, the U.S. money supply totaled about $780 billion.[47] However, at the time, only "$45 million in [...] supernotes [high quality forgeries] [...] have been detected in circulation," and an

---

[44] Lee McIntyre, "Making Money Keeps Getting Easier," *Regional Review*, Quarter 2 (2000).

[45] Irving Fisher, *The Purchasing Power of Money: Its Determination and Relation to Credit, Interest and Crises* (New York: Macmillan, 1911).

[46] Eugene Lerner, "Money, Prices, and Wages in the Confederacy, 1861-1865" *Journal of Political Economy*, 63(1):20–40, 1954.

[47] R. Perl and D. Nanto, "North Korean Counterfeiting of U.S. Currency," 2007, Congressional Research Service report RL33324. Available online at http://opencrs.com/document/RL33324/2007-01-17/.

estimated total of $180 million forged U.S. banknotes were in circulation worldwide. Even using the latter value, forgeries only account for around 0.01% percent of total currency. Under these values, it would take a 200x increase in forged money production to corrupt 1% of the monetary supply of the US and have a mere 1% impact on inflation.

An increase by 200 times seems difficult to achieve. While the production costs of decent quality counterfeiting are decreasing thanks to advancement in cheap printing, scanning, and imaging technologies, high-quality forgeries, such as supernotes, are likely expensive to create. Their fixed costs are significant, because they require equipment as sophisticated as those of national note printing presses. For instance, supernotes must be based on "intaglio" methods from identical presses used by the U.S. Bureau of Engraving and Printing, which are available only from a single manufacturer in Switzerland.[48] In addition, equipment of that type (and distribution channels to inject significant amounts of forged notes into a nation) is not easy to hide. Avoiding police detection implies additional costs.

Finally, while the forgers may be busy trying to increase their production above the current level of $180 million, the U.S. Bureau of Engraving and Printing would still be able to produce "26 million notes in a day with a face value of approximately $907 million"[49]—more than eleven times the estimated number of existing supernotes in circulation.

To achieve such a significant scale in production, a hostile government would have to bear significant costs and engage in large-scale activities that may make it detectable by the U.S. Secret Service agents abroad,[50] or by international organizations such as the Central Bank Counterfeit Deterrence Group (CBCDG). As such, the direct macroeconomic impact of forgeries appears to be negligible.

### D. INDIRECT EFFECTS

However, injecting forgeries into an economy may affect the economy indirectly, as we describe below.

---

[48] Lee McIntyre, "Making Money Keeps Getting Easier," *Regional Review,* Quarter 2 (2000).

[49] From: http://www.bep.treas.gov/document.cfm/18/106. (last visited August 30, 2010).

[50] GAO. Counterfeit U.S. Currency Abroad: Issues and U.S. Deterrence Efforts. Letter Report, 02/26/96, GAO/GGD96-11, 1996.

### 1. MONEY MULTIPLIER EFFECTS

First, with the increasing sophistication of financial markets, the concept of money supply has also kept expanding, leading to possible ricochet effects that can magnify the impact of an inflow of forged money into a nation's monetary base. Economists distinguish M0 (the monetary base that includes all physical currency and the central bank accounts that can be exchanged for physical currency) from M1 (which includes M0 and checking accounts) and M2 (which includes, among others, M1 and most savings and money market accounts). Monetary expansion implies that an increase in one form of monetary instrument—for instance, an increase in banknotes circulated in the economy—will expand through the monetary supply via a "multiplier" effect.[51] As an example of this multiplier effect, a forged $100 dollar bill deposited at a banking institution and then released in circulation results in $200 of fake money being in the system—the $100 deposited, plus the bill that has just been brought back into circulation.[52] Hence, multiplier effects may enlarge the macro impact of mass amounts of forged notes more than a simple contamination of the currency pool would suggest. Historical data on M0, M1 and M2 (available from the St. Louis Federal Reserve[53]) suggest that, since the mid-1980s, $M2/M0$ is between 5 and 8, with smaller values occurring in times of recession, i.e., when the monetary supply is increased. Hence, even considering multiplier effects, the macroeconomic impact of supernotes is unlikely to be very significant.

### 2. POLICY MOTIVATIONS

Solely focusing on the arguably small macroeconomic impact of supernotes, however, misses a more disturbing motivation that may be driving their creation. The forger may not just try to destabilize the target nation's economy, but may instead attempt to inject hard-to-

---

[51] Frederic Mishkin, *The Economics of Money, Banking, and Financial Markets*, Addison-Wesley, 7th edition, 2004.

[52] As noted in the Introduction, "supernotes" are virtually indistinguishable from the legitimate currencies, unless highly specialized and costly forensic analysis is employed.

[53] Federal Reserve Bank of St. Louis. Monetary Aggregates, 2009. http://research.stlouisfed.org/fred2/categories/24.

detect money into a black market in order to enhance other more dangerous operations, such as terrorism.

Since two-thirds of the total U.S. currency supply are held overseas (with the $100 bill being more widespread abroad than in the United States[54]), and since the number of forged notes keeps increasing, while their cost decreases, the risk that increasingly larger amounts will be used to indirectly attack the interests of a nation state is concrete. As discussed above, the US Secret Service is already devoting "a large portion of [its] budget"[55] to anti-counterfeiting activities. Since the U.S. Secret Service 2005 budget request totaled $1.4 billion,[56] even a small fraction of that budget would seem to be more than the actual amount of forged U.S. bills currently in circulation—suggesting strong policy motivations for stronger forms of counterfeit prevention.

### 3. LOCAL DESTABILIZATION EFFECTS

A potentially thorny issue with counterfeit notes is that money travels relatively slowly. A study [57] based on available online data of bill circulation in the United States[58] shows that, after a year, almost a fifth of all notes have traveled less than 50 kilometers. About a quarter have traveled more than 800 km, while the rest (57.3%) have traveled between 50 and 800 kilometers. The study further shows that the distance traveled by currency overall grows with the time since injection of the currency. More precisely, currency does not move outside of a local radius for a while and then "jumps" to a distant place (presumably due to travel), and remains there for a while, before repeating the process.

This relatively low mobility has two very distinct effects from a security perspective. On the one hand, it makes it easier to identify the origin of injection of forged notes, which in turn can lead to easier

---

[54] Lee McIntyre, "Making Money Keeps Getting Easier," *Regional Review*, Quarter 2 (2000).

[55] Lee McIntyre, "Making Money Keeps Getting Easier," *Regional Review*, Quarter 2 (2000).

[56] Evamarie Socha, "Doing Business With the U.S. Secret Service. " *Washington Technology*, 18(24), 2004.

[57] D. Brockmann et al., "The Scaling Laws of Human Travel." *Nature*, 439(2006), 462–465.

[58] *See* http://www.wheresgeorge.com (accessed August 31, 2010).

identification of the perpetrators. On the other hand, it makes it appealing for an attacker to try a massive, localized injection of bills in a highly concentrated market (e.g., New York City) in hopes of destabilizing the local market.

While *monetary* inflation due to an increase in the supply of money in general affects prices at the macroeconomic level, *price* inflation can at times be a local phenomenon (for instance when a local shortage of goods—such as water or food after a natural disaster—causes their prices to rise locally). Consider a scenario where an attacker injects millions in fake currency in Manhattan, e.g., by leaving bags full of bills in Central Park. The effect may be an increase in local prices or a decrease in the trust people naturally afford to legal tender: if the forgeries are indistinguishable from legitimate notes, people will be tempted to try to spend this money, and those receiving the notes, such as merchants, will face the conundrum of accepting all notes, which may cause prices to rise, or refusing certain or all notes for a period of time, thereby causing further disruption.

A variant of the attack would focus on a specific industrial sector, rather than a geographic location, hoping to create ripple effects. For instance, using a massive amount of counterfeit money to buy a commodity like timber could, in the short-term, result in an artificial increase in the price of timber; as a result, this could translate in an artificial increase of the price of products or services associated with timber, for instance increasing construction costs.

### 4. CONSTRAINTS ON THE COST OF COUNTERMEASURES

The above discussion suggests that techniques used to prevent counterfeiting should remain economically efficient to justify changes to the current approach of combining physical security and police intervention. In other words, any (physical) digital cash protocols proposed by cryptographers and computer scientists should reduce the costs of policing without increasing significantly the production and usage costs. In particular, digital solutions to counterfeiting should meet two criteria.

### A. SIMPLE UPGRADE

Any upgrade of the currency design is tightly constrained. Current estimates suggest that the US government spends approximately 10

cents per bill produced,[59] which already represents a considerable increase within the past few years—earlier versions of the Federal Reserve Board reports placed the printing costs at about 5 cents per bill in the early to mid-nineties. Any security extensions to physical bills should impose a negligible overhead over current bill production methods. Techniques that would raise the production cost of a bill to 20 cents, for instance, are unlikely to be adopted.

### B. MINIMAL COST TO THE USERS

A number of failed currency innovations, such as efforts to popularize dollar coins,[60] have shown that people are generally conservative when it comes to currency, and tend to resist drastic changes when they do not perceive any added value. Hence, to gain widespread acceptance, a novel currency design must provide some tangible benefits, yet avoid any possible perceived burdens. Namely, the bill exchange process should not impose any additional transaction cost (monetary or otherwise) to the user, and any verification costs should remain negligible compared to the actual value of a given bill.

## IV. REQUIREMENTS FOR ENHANCED CURRENCY

In this Part, we extend the economic discussion of counterfeiting, and the consideration of economic constraints for currency deployment, to explore the usability and security properties that any enhancement to currency (physical, digital or a combination thereof) should aim for.

### A. USABILITY PROPERTIES

Currency is a universal product, in that almost every individual uses cash. Any changes to currency must therefore preserve the key usability properties of cash. Namely, any proposed solution should satisfy the following usability properties:

---

[59] The Federal Reserve Board 2009, "New currency budget, 2009" http://www.federalreserve.gov/generalinfo/foia/2009newcurrency.htm (last visited Aug. 30, 2010).

[60] While it is still too early to deem the new dollar coin design featuring past presidents a failure, production figures have been decreasing steadily [14], despite the "collector" value provided as an adoption incentive to its users.

**Universal use.** Any enhanced currency should provide the same usage characteristics as current physical cash, offering extreme ruggedness and enabling exchange without digital devices.

**Reusability.** A single bill should be reusable once it is passed from one owner to another. Digital cash, on the other hand, is used only once, and then destroyed.

### B. SECURITY PROPERTIES

To resistant any type of counterfeiting, enhanced currency should fulfill the following security properties:

**Forgery-proof.** Given an electronic verification device, it should be impossible, or at least computationally infeasible, to create a bill that is indistinguishable from one issued by a legitimate entity. In other words, forgers cannot create bills with new denominations or serial numbers; instead, they are limited to high-quality duplication of existing bills.

**Universal verifiability.** We require that bills be verifiable using a commodity electronic verification device. That way, individuals can easily start verifying the correctness of bills. For instance, one of the approaches we consider in this paper is to employ current camera-equipped smart phones as verification devices, since these phones are quickly becoming ubiquitous.

**Useless duplication.** Given an online electronic verification device, it must be impossible to duplicate an existing bill and successfully cash both bills. A single bill has at most a single owner at any given time. This property does not imply that duplicating a bill is impossible, but merely that the duplicated bill should be useless.

**Anonymity.** One of the most salient features of physical cash is anonymity. Even though banknotes do not ensure perfect anonymity,[61] an enhanced currency system should provide a level of anonymity equivalent to that provided by physical cash.

In addition to these usability and security properties, any solution should also guarantee simple upgrades and minimal costs to users, as discussed in the previous Part.

In essence, the above requirements describe the properties that physical cash should ideally satisfy. With the exception of anonymity, current physical cash designs do not satisfy most of the security

---

[61] D. Kügler, "On the anonymity of banknotes," (paper presented at the *4th International Workshop on Privacy Enhancing Technologies (PET'04)*, pages 108–120, Toronto, Canada, May 2004.)

properties outlined above. Digital cash, on the other hand, may not satisfy all usability (and economic) requirements.

Simultaneously meeting all security, usability, and economic requirements is extremely difficult, if not impossible. In the remainder of this paper we contrast several approaches, and show which designs come the closest to satisfying all of our requirements. As pointed out in the introduction, a combination of physical and digital cash, seems more likely to satisfy the economic, security and usability properties than either of physical or digital cash solutions, especially considering consumers' reluctance to completely abandon currency systems with which they are familiar.

## V. PHYSICAL-DIGITAL CASH TECHNIQUES

In this Part we consider a number of techniques for designing a physical-digital cash hybrid, including some novel proposals. We evaluate both the advantages and disadvantages of each system.

### A. BARCODE SIGNATURES

By encoding signatures in 2-D barcodes, we can 1) keep all the properties of existing physical cash, and 2) strengthen the design using cryptographic primitives to make forgery impossible. Simply stated, this technique augments existing bills with an unforgeable cryptographic signature.

**Design.** Since each bill already possesses a unique serial number, $N$, the bill's issuing authority (e.g., federal bank) can sign the serial number and the bill's denomination, $D$, with its private key, $R_{gov}$. The associated public key, $U_{gov}$, should be widely published. While traditional bills only contain $N$ and $D$, physical-digital cash bills contain $(N, D, \{N \| D\}_{R_{gov}})$.

To preserve the ruggedness of physical cash, the digital signature on the bill could be embedded using a 2-D barcode, e.g., PDF417,[62] as shown in Figure 1(a). 2-D barcodes have previously been used for cryptographic verification of metered postage.[63] They allow fast optical scans and are therefore easily verifiable.

---

[62] S. Itkin and J. Martell, "A PDF417 Primer: a Guide to Understanding Second Generation Bar Codes and Portable Data Files" *Technical Report Monograph 8, Symbol Tech.*, April 1992.

[63] J.D. Tygar, Bennet S. Yee, and Nevin Heintze, "Cryptographic Postage Indicia." In *Proc. ASIAN'96*, Singapore, December 1996, pages 378–391.

**Evaluation.** Since the 2-D barcode does not require any electronic circuitry on the bill, the encoded signature will be robust under extreme physical conditions. The encoding process can also employ error-correcting codes to further enhance the robustness of the signature. Thus, barcode signatures satisfy the *universal use* property of physical-digital cash.

As long as the private key $R_{gov}$ is kept secret, and assuming a secure signature scheme, such as RSA[64] or DSA [1], the bills are *forgery-proof*.

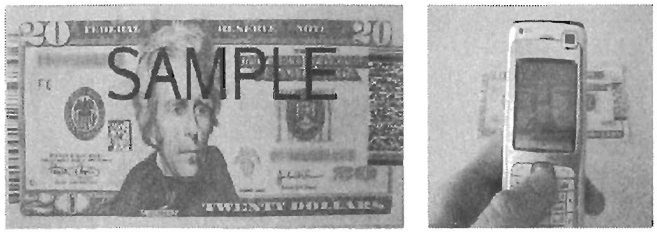


(a) (b)

Figure 1: Barcode Signatures. A sample implementation of physical-digital cash using 2-D barcodes to encode a signature of authentication. Anyone with an appropriate scanner or camera phone can verify that a legitimate institution issued this bill (or one identical to it).

By encoding the signature with a 2-D barcode that can be readily read by commodity camera-based smart phones (as shown in Figure 1(b)), we achieve universal verifiability. In general, a 2-D barcode reader is much simpler than most other verification devices, such as RFID readers. Some smart phones, particularly in Japan and South Korea, are already equipped with barcode reader software. While verification is certainly appealing for higher denominations (e.g., $100 bill), most users would likely not want to verify all bills they have in their possession, and could instead randomly sample the lower denominations bills to verify. In either case, the ability to verify bills for a negligible cost is an important asset.

The manufacturing technology for adding a barcode to a bill is trivial—current bills already contain serial numbers that are printed on each individual bill, and the same technology can also be used to print a barcode. For these reasons, the barcode satisfies the *simple upgrade* property.

---

[64] Ronald L. Rivest et al., "A Method for Obtaining Digital Signatures and Public-key Cryptosystems." *Communications of the ACM*, 21(2) (1978):120–126.

Such a barcode-enhanced bill does not contain more information than a traditional bill: the signature itself can only be used to verify the authenticity of a bill. Thus, the proposed scheme satisfies reusability and anonymity requirements.

However, used alone, signatures cannot enforce the *useless duplication* property. Indeed, a duplicated bill would have the same serial number $N$ as the original (valid) bill, so that $(N, D, \{N \| D\}_{R_{gov}})$ would remain valid. This implies that barcode signatures would not be fully effective solutions against counterfeiting. To achieve the *useless duplication* property, we must turn to additional (or alternate) techniques.

### B. RFID-BASED PROTECTION

An alternative solution, which, as discussed in Part II, was once considered for Euro bills,[65] is to embed RFID chips in bills. Using an RFID chip offers two primary advantages over 2-D barcodes. First, an RFID chip can perform limited computations and can even interact with a reader. Second, while 2-D barcodes are read-only, some RFID chips have writable memory.

**Design.** If we assume the use of tamper-proof RFID chips (we discuss the strength of this assumption below), then one can design a simple protocol, similar to Seeing-is-Believing,[66] to authenticate physical-digital cash. For a bill with serial number $N$, the issuing authority generates a public-private key pair ($K_N$, $K_N^{-1}$), stores $(K_N, K_N^{-1}, \{K_N^{-1}\}_{R_{gov}})$ on the embedded RFID chip, and prints a barcode encoding of $H(\{K_N^{-1}\}_{R_{gov}})$ on the face of the bill, where $H$ is assumed to be a cryptographically secure hash function.

To authenticate a bill, any user with an appropriate reader can transmit a randomly chosen nonce, $\kappa$, to the RFID chip. The chip responds with a signature $\{\kappa\}_{K_N}$ on the nonce, its public key, $K_N^{-1}$, and the certificate, $\{K_N^{-1}\}_{R_{gov}}$, for its public key. The reader checks the signature using the public key provided and checks that the hash of the certificate matches the commitment printed on the face of the bill.

**Evaluation.** RFID chips will be less tolerant of daily wear and tear and extreme environmental conditions than the original bill. As

---

[65] Junko Yoshida, "Euro bank notes to embed RFID chips by 2005," *EE Times* (2001). http://www.eetimes.com/ story/OEG20011219S0016, December 2001.

[66] Jonathan M. McCune, Adrian Perrig, and Michael K. Reiter, "Seeing-Is-Believing: Using Camera Phones for Human-Verifiable Authentication." In *Proc. IEEE Security and Privacy*, May 2005.

such, an RFID-based approach may not fully satisfy the *universal use* requirement. Further, at present, an RFID approach does not satisfy the *universal verifiability* requirement, as RFID readers have not yet penetrated the consumer market. Likewise, embedding a computational device in each bill would significantly raise the cost per bill (up to $1, that is, a 10-fold increase [67]) and alter production methods. While improvements in RFID technology may remedy this drawback, this technique currently does not provide a *simple upgrade*.

Since the data stored on the RFID chip does not include any information about the owner of a bill, this technique achieves both *reusability* and *anonymity*. A perfectly secure RFID chip may make forgery and duplication impossible, thereby directly enforcing the desired *forgery-proof* and *useless duplication* properties. Unfortunately, trusting the security of an RFID chip is an extremely strong assumption, as has been evidenced by existing attacks.[68] It remains an open question whether similar techniques can be developed using insecure RFID chips.

Finally, another disadvantage of RFID chips is that they can be read remotely, potentially enabling a thief to determine the amount of money a potential victim is carrying. Similar to the vulnerabilities of the new RFID-based US passport,[69] adding RFID tags to bills would raise numerous new vulnerabilities.

### C. PHYSICAL ONE-WAY FUNCTIONS

A different way to ensure the useless duplication property is to embed a physical one-way function in each bill.

**Design.** Physical one-way functions can be implemented, for instance, by randomly sprinkling bits of optical fiber in the fabric of each banknote,[70] or by using magnetic polymers.[71] Each bill has

---

[67] Junko Yoshida, "Euro bank notes to embed RFID chips by 2005." *EE Times* (2001), http://www.eetimes.com/ story/OEG20011219S0016, December 2001.

[68] S. Bono et al., "Security Analysis of a Cryptographically-Enabled RFID Device" (paper presented at USENIX Security, Baltimore, MD, August 2005).

[69] Bruce Schneier, "Renew Your Passport Now!" Last modified October 2006, http://www.schneier.com/crypto-gram-0610.html.

[70] G.J. Simmons, "Identification of Data, Devices, Documents and Individuals." In *Proc.* Presented at the 25th Ann. Intern. Carnahan Conference on Security Technology, Taipei, Taiwan, ROC, October 1991, pages 197–218.

unique characteristics due to the length and orientation of the fiber strands or polymer present in its fabric, and it is extremely hard to produce a copy of the bill with an identical physical configuration.

Exposing the bill to a light (or magnetic) source under different conditions (e.g., different angles) yields a unique characterization of the structure of the bill, this can be numerically encoded and printed on the bill. Verification is a matter of exposing the bill to the same conditions and matching the information printed on the bill. Combining this scheme with a signature scheme, e.g., by signing the value characterizing the physical structure of the bill can further ensure the forgery-proof property.

**Evaluation.** This approach has the merit of providing enhanced security without changing the way people would use bills. Three important open problems remain, however, regardless of the physical one-way function used. First, the manufacturing cost of such bills is hard to assess, but is certainly much higher than the current production cost. Second, fibers, or polymers may break or get dirtied easily, resulting in genuine bills failing the verification process. Third, the equipment needed to verify such enhanced bills is likely to be too high an investment for most merchants, let alone individual users. Due to the cost of the required verification equipment, forgeries may travel undetected in the monetary network for considerable amounts of time.

As such, physical one-way functions do not easily satisfy *universal verifiability*, *simple upgrade*, or *universal use*. However, as we discuss later, we believe physical one-way functions may be very useful when deployed in conjunction with other techniques.

### D. Paper Fingerprinting

An idea closely related to physical one-way functions is discussed in Clarkson et al.,[72] which shows that imperfections in paper can be measured, and used to characterize the uniqueness of a piece of paper.

**Design.** Clarkson et al. show that, for typical copy paper, one can measure imperfections in the grain and structure of the paper. The technique proposed consists in scanning the paper at different angles with (cheap) commodity scanners, and to combine these

---

[71] H. Hoshino et al., "Object to be checked for authenticity and a method for manufacturing the same," February 1997. US Patent nr. 5,601,931.

[72] William Clarkson et al., "Fingerprinting Blank Paper Using Commodity Scanners" (paper presented at the IEEE Symposium on Security and Privacy, May 2009).

measurements into a hash value. The hash can then be embedded on the paper as a 2-D barcode, similar to the barcodes we discussed earlier for authentication. Thus, one could modify the 2-D barcode we proposed earlier to include $(N, D, \{N \| D \| S\}_{R_{gov}})$ where $S$ is a hash characterizing the structure of the paper. Doing so, we can bind a given sequence number and denomination to a paper structure, which implements *useless duplication*. Indeed, a duplicate would have a different paper structure (due to paper imperfections being random), and consequently a different value $S$, which would not match the printed barcode.

**Evaluation.** The idea of fingerprinting paper makes *universal verifiability*, *simple upgrade* and *universal use* considerably easier to satisfy than the physical one-way functions discussed before. Note however, that a cellular phone reading the 2-D barcode would need to obtain $S$ by other means (e.g., external input coming from a scanner), which would make the verification of $S$ relatively cumbersome in general.

Furthermore, paper fingerprinting needs to be more thoroughly evaluated before we can be convinced of its feasibility for preventing monetary counterfeits. First, currency is printed on specialized paper, made for instance of cotton or plastic fibers and determining to what extent existing techniques apply to currency paper is an open question. Second, currency is subject to considerably rougher handling than typical documents. Clarkson et al. show the resiliency of their method to printing, scribbling and wetting and drying, which seems very promising. However, compared to typical printed documents, bills are smaller, constantly folded (which may alter the structure significantly), and frequently dirtied. Overall, we would expect the fingerprinting to be considerably more complex and to potentially lead to false positives.

With these caveats in mind, the idea of fingerprinting paper, much like the idea of using physical one-way functions, seems extremely appealing when combined with the other techniques we describe in this paper.

### E. CENTRALIZED VERIFICATION

Both centralized and decentralized verification (discussed in Part V.F) attempt to achieve the *useless duplication* property. While neither provides a completely satisfactory solution, both represent interesting points in the design space.

**Design.** One simple way of making duplication more costly for counterfeiters is to keep a database of issued serial numbers at the

issuing central bank, and require that all banks are able to quickly verify whether a given serial number has already been deposited or not. We can thus ensure that two bills with the same serial number cannot be deposited at the same time. Adding a cryptographic signature on the bill would both prevent the introduction of illegitimate serial numbers and detect the duplication legitimate serial numbers. Without the cryptographic signature, this technique directly applies to unmodified physical cash, but it offers weaker properties, since it can only detect the introduction of illegitimate serial numbers when the bills are deposited at a bank.

**Evaluation.** Given that centralized verification utilizes unmodified physical cash, it clearly meets our *universal use* and *reusability* goals. It imposes no additional production (marginal) costs, making it a *simple upgrade* to the printing process. However, it does impose fixed costs on the central bank, which must maintain the serial number database, as well as on the member banks that must constantly monitor and report on the serial numbers entering and leaving their control. Centralized verification minimally impacts the traditional *anonymity* of physical cash, since the bills remained unchanged, and serial number data is already available at the member banks.

Without barcode signatures, centralized verification of serial numbers is only partially *forgery-proof* and provides only limited verifiability, since only banks can perform the verification procedure. Further, duplicate bills can remain in circulation undetected for extended periods of time. In fact, until one of the bills is deposited, not even the central bank knows that duplication has occurred.

### F. DECENTRALIZED VERIFICATION

Ideally, with a distributed verification scheme, we could achieve instant detection of duplicates, such that no one would accept a duplicate bill. Decentralized verification attempts to achieve this property by enabling individuals and merchants to perform real-time validation of bills they receive. The novel system we discuss here offers stronger properties, but it also imposes larger costs and may introduce new vulnerabilities. While it does not offer a perfect solution, it does suggest a direction for further research. Indeed, the increasing ability of engineers to design large-scale, distributed databases may prove a valuable asset in counterfeit prevention. Relying on the assistance of online servers, which would have been unthinkable only a few years ago, is becoming a credible proposition.

**Design.** At a high level, a decentralized database (perhaps hosted by various member banks or other governmental agencies) associates each bill's serial number with a cryptographic "lock bit". Once a bill is locked, only the current "owner" of the bill can unlock it. To transfer ownership of a locked bill, the current owner cryptographically unlocks it and allows the new owner to lock it. Participants can check the current state of a particular bill's lock bit and may refuse to accept a locked bill.

While we refer to this bit as a "lock bit" throughout, it merely implements a warning system, rather than an actual enforcement (blocking) mechanism: *users can elect to receive bills even if the lock bit is on*, but do so at their own risk.

Furthermore, dealing with legacy users (i.e., those that cannot check a bill's lock status) requires additional precautions. In general, before transferring a locked bill to a legacy user, the current owner should unlock it so that the legacy user can make use of it unhindered. By default, all bills dispensed by an ATM to a legacy user would be unlocked (or locked with a null value) by the issuing bank. Participating users would then take ownership of the bills by immediately locking them.

On a related note, since a legacy user cannot check the status of a bill's lock bit, a participating user might accidentally or maliciously provide them with a locked bill. A similar problem arises if a participating user loses the cryptographic material necessary to unlock their own bills. To address this problem, the decentralized verification service must be backed by the central bank. We assume that the central bank can distinguish a duplicate from a real bill through some, possibly costly, verification process. For instance, physical one-way functions or the type of paper fingerprinting described above could assist in the bank's verification process. Indeed, used as a back-up verification system, physical one-way functions do not need to have the same level of robustness as when used as the primary mechanism to prevent duplication.

With this decentralized verification system in place, a user could deposit a locked bill at a bank in a procedure similar to that used for checks today. The bank would send the locked bill back to the treasury to verify its authenticity. If the bill is authentic, the bank will credit the value of the bill to the user's account, regardless of its lock status.

**Implementation.** In Appendix A, we describe in more detail how such a distributed locking scheme can be implemented in an anonymous manner, using one-time public/private key pairs. The scheme essentially consists of a short series of messages between two participants in a transaction and the bank.

**Evaluation.** Given that the only modification of the actual physical currency is the encoding of each bill's serial number in a machine-readable form, decentralized verification achieves the same strong *universal use* property as the barcode signatures, and as far as the production process is concerned, only requires a *simple upgrade*. While transfers between participants become more complicated than with standard physical cash, physical-digital cash with decentralized verification can still be used by, and exchanged with, legacy users that do not have the appropriate electronic devices. This also implies that this technique satisfies the *reusability* requirement.

Both the locking procedure described above (and any checks on the lock status) will fail if the serial number provided does not exist, so anyone with a scanner can determine the authenticity of a particular bill, making the currency *forgery-proof*. Since anyone with an online connection can query the lock status of a particular bill, this technique also provides *universal verifiability*. Current smart phones have access to a high-speed Internet network enabling them to establish a secure communication channel with the bank. Short-range wireless communication capabilities can be secured using known techniques[73], and used to transfer bills between participants.

The stored information for each bill consists of the double $(N, \lambda)$. With about 20 billion bills currently in circulation,[74] and the conservative assumption that each double $(N, \lambda)$ requires 64 bytes, the total size of the database is about 1 TB, a small number compared to other existing highly-available databases such as web indexes[75].

Decentralized verification provides a reasonable level of protection against duplication by using a distributed network of verifiers to enforce the principal of *useless duplication*. A participant in the system that receives an unlocked duplicate should immediately lock it, preventing any of the copies from being locked by other participants. Transferring a duplicate to another participant has a similar effect. If a forgery does occur, it drives all bills back to the bank, since merchants will not accept duplicates of a bill once the first bill has been locked. This allows easier monitoring and can yield clues for enforcement.

---

[73] Jonathan M. McCune, Adrian Perrig, and Michael K. Reiter, "Seeing-Is-Believing: Using Camera Phones for Human-Verifiable Authentication," In *Proc. IEEE Security and Privacy*, May 2005.

[74] U.S. Department of Treasury. Treasury Bulletin, June 2007. http://www.fms.treas.gov/bulletin/.

[75] Fay Chang et al., "Bigtable: A Distributed Storage System for Structured Data" (paper presented at ACM/USENIX OSDI'06, , 2006).

The primitive does come at a usage cost, though: users forgetting to lock unlocked bills may actually see their valid currency be subsequently rejected, *if forgeries with the same number are present in the system.* Given the relatively modest amount of high quality forgeries in circulation, compared to the total number of bills, we think that this is unlikely to occur. Recall that our threat model assumes a relatively limited number of very high quality forgeries, rather than a vast amount of low quality forgeries, which can be detected much more easily.

Because one can implement the exchange protocol using only transient random numbers that cannot be matched to any real-world identity, the transfer protocol does not in itself appear to degrade *anonymity* of physical cash. A thornier issue is that of accesses to the online database. In the exchange protocol we propose, the bank *B* knows when user *A* wants to spend the bill *N*, since *A* contacts *B* directly. By extension, as long as the bills are passed between principals that use bill scanners and locking primitives, *B* has a way of reconstructing the whole transaction chain. Because the communications between *A* and *B* never involve the names of the principals (no message include the names *A* or *C*), the problem can be solved by using anonymous communication primitives[76] that make it impossible for the bank to identify *A*. This system could achieve reasonable levels of *anonymity*, possibly at the expense of added latency.

Decentralized verification, thanks to the (un)locking primitives, can also help combat theft. A wallet full of locked bills is useless to a thief. Ownership has not been relinquished, and the money cannot be deposited or exchanged with any participant in the system. Also, the owner of the locked bills retains the serial numbers and unlocking codes for the stolen bills, and can provide this information to the authorities: The thief cannot deposit the money at a bank by claiming to have lost the unlocking codes. These benefits may encourage adoption, since only participants in the system will have this protection. As a drawback, under that scheme, legitimate users forgetting to request unlocked bills when they have a right to do so could have a harder time justifying the money is indeed theirs, but we presume this type of user error would become rare once people get more familiar with the system.

---

[76] E.g., David Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms" *Comm. ACM*, 42 (1981); Roger Dingledine, Nick Mathewson, and Paul Syverson, "Tor: The Second-Generation Onion Router" (paper presented at the 13th USENIX Security Symposium, August 2004).

In sum, none of the techniques discussed here perfectly meet all requirements outlined in Part IV. However, a combination of these techniques represents interesting and useful building blocks for future physical-digital cash schemes.

## VI. SECURITY ANALYSIS

The various techniques outlined above for implementing physical-digital cash raise a number of questions regarding possible vulnerabilities of physical-digital cash.

### A. COMPROMISED PRIVATE KEYS

If the private key $R_{gov}$ used for signing the bills is compromised, physical-digital cash is not forgery-proof anymore, and the security level degrades to that of physical cash. Unfortunately, the public may rely on the cryptography as hard evidence that a bill is legitimate, rather than also checking other security signs, such as physical watermarks.

While the issuing government should immediately replace the key pair $R_{gov}$, $U_{gov}$, recalling all bills signed with the compromised key may prove problematic. Massive recalls have been shown possible in practice, e.g., by the recent shift from all national European currencies to the Euro, but large-scale recalls are costly and take several years to be effective. A possible way to mitigate the risk of a key compromise is to use different keys for each denomination, e.g., $20 bills, produced at a given facility, and with a limited lifetime. Limiting the number of bills involved would facilitate a relatively rapid recall in case of a key compromise.

### B. FAKE SIGNATURES

Another class of attack consists of attacks on the signature itself. We are not concerned by cryptographic attacks here, but by physical attacks on the signature information. For instance, fake bills may be produced with missing or incorrect digital signatures. A missing signature is very easy to notice, but while an incorrect signature can be easily detected using a bill scanner, it is not easy to detect in the off-line realm: there is no obvious visual distinction between a good and a bad signature.

Worse, the visible presence of a digital signature (e.g., the presence of a 2-D barcode) may convince users that the bill is good, even in the absence of verification. From a psychological standpoint, a

bill may look more trustworthy just because of the apparent presence of a digital signature, even though other physical indicators, e.g., the quality of the paper, or the presence of a watermark, may be questionable.

### C. ROGUE FINANCIAL INSTITUTIONS

Serious problems may arise when a rogue financial institution (e.g., bank, foreign currency exchange shop) participates in exchanges. One whole class of attacks can be characterized as "money laundering," that is, in the context of counterfeit money, exchanging fake bills for good bills. The simplest instance of such an attack is when a dishonest merchant tries to pass bad bills onto customers. This type of attack is not new, and in fact already affects the existing physical cash network. Countermeasures are simple: in the physical cash network, individuals are supposed to check the physical properties of a given bill. In the physical-digital cash network, individuals can use readers (e.g., applications on their smart phones for barcode signatures, miniaturized RFID readers, etc., depending on the technique employed) to thwart this problem.

A more elaborate version of money laundering involves an attacker colluding with a rogue bank, which cashes counterfeited bills produced by the attacker without checking them. The counterfeited bills are then sent to the currency exchange office of the bank, where they are exchanged for good foreign currency bills from unsuspecting tourists. As long as bills are not verified and no one attempts to lock them, they may travel in the network. Monitoring banks is a plausible countermeasure against such an attack. Compared to the large number of bill users, there are relatively few banks in the world, so a centralized authority (e.g., a treasury department) could monitor them effectively. Recent events indicate that such monitoring already exists in practice.[77]

Another variant on the money laundering scheme is that used by a rogue foreign exchange shop that does not just accept, but also gives out popular foreign currency (e.g., U.S. dollars) in a different country (e.g., Japan). These shops are much less regulated and less controllable than banks. However, for a popular currency, we expect the flow of money to be mostly from the tourists to the foreign exchange shops (e.g., backpackers exchanging US dollars for local

---

[77] S. Mihm, "No Ordinary Counterfeit," *New York Times Magazine*, July 23, 2006, 36.

currency), so the impact of this attack should be limited. Further, in all money laundering attacks, counterfeit bills are detected as soon as the bill is deposited at a legitimate institution, or passed to an individual equipped with a bill scanner.

### D. DENIAL OF SERVICE

While currency is usually not subject to denial of service attacks, an *online* locking mechanism as described in Part V. does potentially introduce new DoS vulnerabilities. Assume that the adversary can create duplicates of existing bills at will. For a nation-scale adversary, this can be done relatively easily, for instance, by asking a large number of people to take pictures of valid bills, or to have a few spies take pictures of a large number of bills stored in banks. Now, consider one legitimate note $L$ with serial number $N$, and its copy $F$, which has the same serial number $N$. $L$ is unlocked as soon as it is passed from a merchant, bank, or individual with the proper equipment to a "legacy principal" which does not have any means to lock bills. The attacker can figure out if $L$ is unlocked by repeatedly trying to lock the note using a null value as the current locking value. As soon as the note $L$ is detected to be unlocked, the attacker issues $F$. If $F$ is locked before $L$, $L$ becomes impossible to spend *even though it is a valid bill.* The only way for the unfortunate owner of $L$ to get his money is to confirm with the treasury that $L$ is, in fact, a valid bill, relying on physical features of the bill, e.g., a physical one-way function.

The central bank may then decide to recall the serial number $N$, but this gives the attacker a way of destroying money, which can lead to sabotage operations. For instance, the attacker may start issuing many copies of bills to disrupt the monetary system by having a large number of users requesting that the treasury check their bills, and having, as a final result, vast amounts of serial numbers destroyed. While the attacker does not gain any money from such a destructive scheme, this type of attack may exert significant pressure on the monetary system targeted.

While potentially serious, these vulnerabilities already exist with physical cash. The presence of a verification system does improve the situation, by making it easier and faster to detect criminal activity. Although the issue of locking a bill held by a legacy principal seems cumbersome at first glance, since the principal will need to deposit the bill at a bank for verification, this action is always due to criminal activity. This should be fairly infrequent, and actually provides an incentive for people to adopt verification devices.

## VII. DISCUSSION

With the objective to significantly strengthen current bills against high-quality monetary forgery, we have highlighted a set of requirements that are needed for a viable solution. Then, we have looked at possible ways to implement these requirements, by augmenting bills with cryptographic material directly embedded in the bill. We have considered optically verifiable cryptographic signatures expressed as 2-D barcodes, RFID chips, physical one-way functions, centralized verification and decentralized verification.

None of the techniques we investigate or propose, when used in isolation, satisfies all the properties we would like to enforce. Each is characterized by unique trade-offs. For some of them, the usability or upgrade costs associated with their implementation may outweigh the current expected benefits of such an implementation. However, a combination of these techniques – for instance, coupling a decentralized verification protocol with optical signatures, and with physical one-way functions serving as back-up – could come very close to implementing all the security and usability requirements we described, and certain economic conditions may justify their adoption.

To avoid deployment issues, decentralized verification schemes should be designed to accommodate legacy users who do not wish to participate in the online verification scheme. More importantly, deployment need not be universal. By driving forgeries back to the banks quickly, a decentralized system should work effectively as a deterrent against counterfeiting, even in the absence of wide deployment. Likewise, it is also possible that implementing only a subset of the techniques discussed in the paper may be enough to discourage most fraud. A design solely based on 2-D barcodes will limit forgeries to duplication of existing bills, and even such duplication would be readily detected.

From an economic standpoint, the added costs of a basic scheme, solely consisting of embedded barcodes, are extremely low: bar-coded signatures could be added to existing notes using available presses. Its impact on the attacker's cost benefit analysis would nevertheless be significant as an attacker would now be physically limited in his ability to create fake notes by the number of legitimate notes he can put his hands on. This would increase his costs and reduce the amount of notes it can produce in a given amount of time. The attacker will still be able, of course, to make many copies of the same note and attempt to flood a market with that note. Our economic discussion suggests that such a simple solution may actually be the most desirable given

the relatively small impact of forgeries on macroeconomics, and the reasonably high efficiency of policing methods already in place.

If the economics of counterfeiting were to change – for instance, if we start noticing large cities being flooded with rogue currency – it may make sense to consider combining offline primitives with a verification system. Such a strategy would decrease the cost of detecting duplicates for the defender. Indeed, it would make it more likely that the counterfeit note will be detected, more likely that it will be detected early in its introduction into the system, and more likely that its exchanges will be tracked.

The added costs of a complete scheme combining physical protection with decentralized verification would be significant. Such a scheme indeed implies higher fixed costs, including equipment costs and transition costs, such as the cost of adopting scanners for banks, merchants, and, in the long term, principals, as well as higher variable costs (e.g., the transaction costs associated with the time spent scanning currency and connecting to an online database whenever a note must be locked or unlocked).

Nevertheless, similar fixed transition costs have been incurred before by large economies – for instance, the transition to electronic check-out cashiers, or the transition to the Euro within some EU member states. One advantage of the scheme we have discussed is that it allows for an arbitrarily long transition period, since the decentralized verification system is essentially used as a warning tool; only when bills get back to the bank does detection (and possibly destruction) of forged banknotes take place. Overall, the variable transaction costs would chiefly depend on how technology will blend into the everyday usage of cash and how seamlessly the locking and unlocking process can be integrated into existing merchant infrastructures. The economic discussion we present in this paper tends to suggest that, under the current conditions, the costs would actually outweigh the benefits of such an implementation; however, we also discuss a number of plausible scenarios/attacks that could considerably alter the economic proposition.

More generally, a deeper consideration of the economics at stake in the production and deployment process of counterfeit-resistant bills warrants further research. We hope that our initial approaches will encourage additional efforts in this important area.

## AN ONLINE VERIFICATION IMPLEMENTATION

To implement the online verification scheme described in Part V.F, the "bank" (e.g., the central bank or the treasury), denoted *B*,

maintains a distributed database that contains an entry for each bill in circulation. Each entry is of the form $(N, \lambda)$, where $N$ represents the bill's serial number and $\lambda$ indicates the lock status of that bill. If $\lambda = \varnothing$, the bill is unlocked, whereas any non-zero value indicates that it is locked. To facilitate the automation of the steps described below, each bill's serial number should be encoded in a machine-readable form such as a 2-D barcode.

To lock an unlocked bill with serial number $N$, a principal (e.g., an individual or merchant) $A$ picks a random value $\mu_A$ and computes $\lambda_A = H(\mu_A)$, where $H$ is a one-way hash function assumed to be secure, i.e., at least weak-collision resistant. Using the bank's public key,[78] $A$ securely transmits $(N, \lambda_A)$ to the bank. The bank will update the database appropriately. We summarize these steps below:

1. $A \rightarrow B: \{N, \mu_A, \varnothing\}_{U_{gov}}$

2. Retrieve $(N,\lambda)$, check $\lambda=\varnothing$, store $(N, \lambda_A)$.

To transfer the bill to another principal, $C$, $A$ will unlock the bill and simultaneously lock it under $C$'s lock value. To simplify the presentation, assume $A$ and $C$ have established a secret key $K_{AC}$, and let denote the authenticated encryption of a message $M$. When the transaction is about to take place, $C$ picks a secret random value $\mu_C$, and computes its hash $\lambda_C = H(\mu_C)$. The following bill transfer protocol takes place:

1. $C \rightarrow A: \{\lambda_C\}_{K_{AC}}$

2. $A \rightarrow B: \{N, \mu_A, \lambda_C\}_{U_{gov}}$

3. $B$: Retrieve $(N, \lambda_A)$, check $\lambda_A=H(\mu_A)$, store $(N, \lambda_C)$

4. $B \rightarrow A: \{N, \lambda_C\}_{R_{gov}}$

5. $A \rightarrow C: \{N, \lambda_C\}_{R_{gov}}$

---

[78] As before, the bank's public key is $U_{gov}$ and its private key is $R_{gov}$. These keys need not be identical to the keys used to authenticate bills through the 2-D barcode. The bank's signature on message $M$ is given by $\{M\}_{R_{gov}}$, and public-key encryption of $M$ is denoted by $\{M\}_{U_{gov}}$.

That is, $C$ gives $A$ the lock value $\lambda_C$, which $A$ forwards to the bank along with her unlocking value $\mu_A$. The bank replaces $\lambda_A$ with $\lambda_C$, effectively updating the "owner" of the bill, before communicating the change back to $A$. Finally, $A$ relays this information to $C$, proving that the lock value has been updated, and physically transmits the bill to $C$.

The key feature of this scheme is that, if the values $\mu_A$ and $\mu_C$ are truly chosen at random, bills can be locked to a given individual without making this individual traceable. Basically, ($\mu_A$, $\lambda_A$) and ($\mu_C$, $\lambda_C$) are used as one-time public-private key pairs.

The above exchange protocol assumes that both $A$ and $C$ are able to participate in an online exchange. If $C$, for example, is unable to participate in an online exchange, because it does not have a bill scanner or does not wish to use it, then $A$ simply unlocks the bill and leaves it in the unlocked state. This can be accomplished with a protocol similar to the locking protocol, namely:

1. $A \rightarrow B : \{N, \mu_A, \varnothing\}_{U_{gov}}$

2. $B$ : Retrieve $(N, \lambda)$, check $\lambda = H(\mu_A)$, store $(N, \varnothing)$.