

The Postmodern Ponzi Scheme: Empirical Analysis of High-Yield Investment Programs

Tyler Moore¹, Jie Han² and Richard Clayton³

¹ Computer Science Department, Wellesley College, USA tmoore@cs.wellesley.edu

² Computer Science Department, Wellesley College, USA jhan@wellesley.edu

³ Computer Laboratory, University of Cambridge, UK
richard.clayton@cl.cam.ac.uk

Abstract. A High Yield Investment Program (HYIP) is an online Ponzi scheme, a financial fraud that pays outrageous levels of interest using money from new investors. We call this fraud ‘postmodern’ in that sophisticated investors understand the fraud, but hope to profit by joining early. These investors support ‘aggregators’ – reputation websites that track the status of HYIPs. We examine 9 months of aggregator data and show that there is no evidence of collusion between different aggregators. We use their data to measure the time until HYIPs collapse, finding – perhaps unsurprisingly – that longer lifetimes are associated with lower interest payments and longer mandatory investment terms. We look at the role of digital currencies in supporting HYIPs, finding that a handful of systems dominate. Finally, we estimate that this type of criminality is turning over at least \$6 million/month and set out ways in which it might be disrupted.

1 Introduction

A High Yield Investment Program (HYIP) is an online version of a financial scam in which investors are promised extremely high rates of return on their investments. Payments are made to existing investors from the funds deposited by newcomers, continuing until insufficient funds remain and the scheme collapses. Similar schemes have operated in the offline world for 150 years or more and are often called *Ponzi schemes* after a famous swindler in 1920’s Boston.

Despite being illegal to operate in most jurisdictions, there are a considerable number of active HYIP websites at any given time. We call them ‘postmodern’ Ponzi schemes because we believe that many of the investors are well aware of the fraudulent nature of the sites, but are of the opinion that by investing at an early stage – and withdrawing their money before the scheme’s collapse – they will be able to make a profit at the expense of less savvy investors.

An extensive online ecosystem has developed in support of HYIPs, involving discussion websites, digital currencies, and third-party ‘aggregator’ websites that track HYIP performance. These aggregators list dozens of active HYIPs, tracking core features such as interest rates, minimum investment terms and funding options. They operate forums in which individuals can report their experiences;

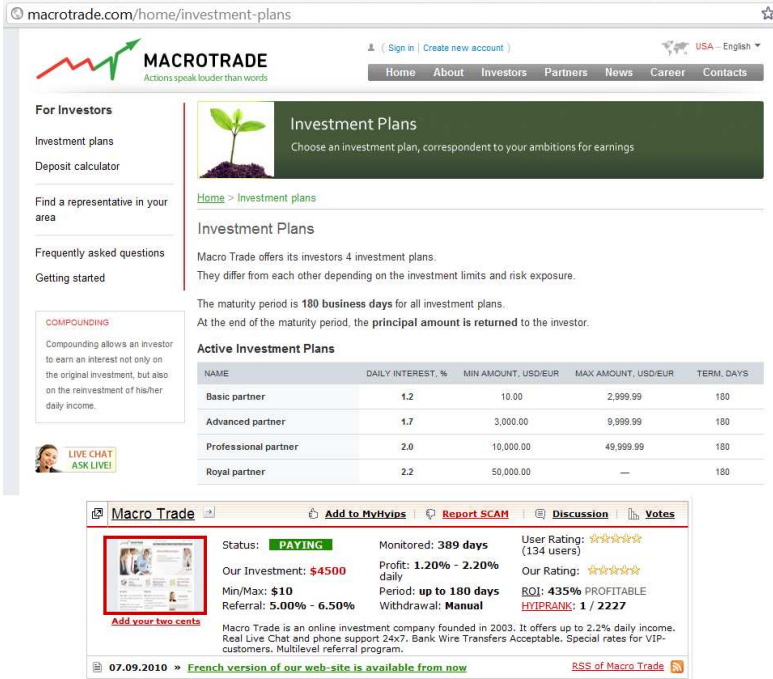


Fig. 1. Screenshot of HYIP macrotrade.com and the corresponding entry on the aggregator hyip.com.

but more significantly, the aggregators appear to make their own investments in some of the HYIPs and report on when interest payments cease. As an illustrative example, Figure 1 shows a screenshot of the HYIP macrotrade.com, along with its entry on the aggregator website hyip.com.

We have spent many months collecting data from HYIP websites and aggregators to measure the extent of HYIP activity, so that we can improve our understanding of this particular type of online criminality.

In Section 2 we explain our data collection and measurement methodology. In Section 3 we discuss our evidence as to whether the aggregator sites are making truthful¹ reports. In Section 4 we examine HYIP lifetimes and investigate the extent to which it is possible to predict their collapse. In Section 5 we discuss the role of ‘digital currencies’ in this ecosystem and then in Section 6 we estimate the scale of this particular type of online criminality and discuss various ways that it might be discouraged, if not entirely stamped out. In Section 7 we survey related work and finally in Section 8 we summarize what we have learned so far and consider what further work might reveal.

¹ We avoid the word ‘honest’ because this is not an appropriate word to use in conjunction with criminal activity.

2 Data Collection Methodology

We term the websites that provide reputation services for HYIP programs ‘aggregators’. Given that HYIPs are confidence games that keep growing so long as new investors can be recruited, these ratings are potentially very powerful indicators of HYIP success or failure. From a Google search for “HYIP” issued in November 2010, we identified 9 aggregator websites to monitor (myhyip.com, maxhyip.com, iehyip.com, hyipranks.com, hyipmonitor.com, hyipinvestment.com, hyip.com, hothyips.com, and everyhyip.com).

Between November 17, 2010 and August 21, 2011 we made daily visits to each aggregator website (with the exception of four days in November and December 2010 due to a bug in our crawler). We parsed the pages we fetched to extract the key characteristics of the HYIPs they listed: interest rate, investment term(s), user and aggregator ratings, along with payment status (i.e., paying, not paying). The names aggregators use for each of these fields varied slightly, so we manually unified the terminology and stored each observation in a database. A total of 141 014 observations were made.

All the aggregator websites provide links to the HYIPs, though some of these links pass via an interstitial page. From January 2011 onwards, we determined the URL of each of the HYIPs and captured the WHOIS record for each HYIP domain. Our automated system also visited each HYIP website, and stored the source files linked to or loaded from the home page. These daily visits to the HYIP websites were made over Tor²; its anonymity properties help ensure that the website would not be able to identify us or trivially connect our visits.

2.1 Measuring HYIP activity

We have used the collected data to derive several key measurements, whose calculation we now describe.

Linking HYIP records across aggregators. Unfortunately, it can be difficult to determine when two aggregators are reporting on the same HYIP. We use the website address of the HYIP as a canonical identifier, but when we failed to ascertain this (e.g., the HYIP website was shut down before we followed the link), we have compared the names that the aggregator gave to HYIPs – stripping out whitespace and punctuation and doing a caseless match.

The 9 aggregators listed 1 576 distinct HYIPs – of these, 211 did not resolve to a website and could not be identified as an HYIP which had ever been resolved. 595 HYIPs appeared on more than one aggregator website, while the other 981 appeared only once. It is likely that some of the 981 unique HYIPs are duplicates that we failed to link up; however, we treat them as distinct in our study.

² <http://www.torproject.org/>

Measuring HYIP lifetimes. One key measure of HYIP performance is how long after initial creation the scheme collapses. Identifying when a website is ready for business is impracticable, so we deem the HYIP lifetime to be the elapsed time between the HYIP’s first appearance as reported by an aggregator site (which we believe will be contemporaneous with the first accounts being created) and its eventual disappearance from that aggregator (invariably because the HYIP has collapsed and is no longer paying).

Normalizing profit rates, investment terms, and expected payouts. There is enormous variation in the interest rates promised by HYIPs, from the outrageous 440% in 10 minutes offered by `top-capital.com` to the comparatively modest 1–2% per day offered by `macrotrade.com`. Many HYIPs offer a menu of investment choices that vary by investment level and term, just as a legitimate bank does for their certificates of deposit (CDs).

For this paper, we start by normalizing the published interest rates and investment terms to a daily rate. We then compute an expected payout value that is standardized across HYIPs. To arrive at the expected payout, we had to infer a model of how investments grow over time. Subtly different phrasing must be interpreted differently, as indicated in the following table:

Phrase	Interest Rate	Investment Term	Expected Payout
$x\%$ for y days	$x\%$	y	$x \times y \times$ principal
$x\%$ in y days	$\frac{x}{y}\%$	y	$x \times$ principal
$x\%$ after y days	$\frac{x}{y}\%$	y	$x \times$ principal
$x\%$ daily	$x\%$	-	-

In every case we do *not* compound daily on the current value, but compound on the original principal. In other words, we do not assume that any of the interest that is paid out will be reinvested. We take this approach because it is consistent with the returns on investment (ROI) reported by the vast majority of aggregators. Additionally, if HYIP investors are indeed ‘postmodern’ and know to take profits as rapidly as possible, then their strategy will be to be put in their maximum investment at the earliest possible time and never add to it.

3 Can the Reports of HYIP Aggregators be Trusted?

Given that all HYIPs are fraudulent, it is natural to ask whether the reports from aggregators should be trusted. While ascertaining ground truth is impossible, we have devised a number of measurements to assess the relative accuracy of data reported on HYIPs.

In particular, 595 of the 1576 HYIPs (38%) are tracked by at least two aggregators and so we can compare the reports about the same HYIP across different aggregators. If there is rough consensus then, either the aggregation sites are in a universal conspiracy, or they are independently assessing the HYIPs in a truthful manner.

3.1 Reporting of HYIP Attributes

We determined what the aggregators reported to be the maximum and minimum investment levels allowed by the HYIP, the referral rates offered to affiliates for signing up new investors, and the withdrawal method offered (automatic, manual, or instant). When we collate this information and look for similarity we get these results:

	Investment		Referral Rate		Withdrawal
	max	min	high	low	type
Perfect Agreement	0.40	0.87	0.44	0.43	-
Diversity Index	0.72	0.94	0.77	0.75	0.88

The first row of this table reports the fraction of HYIPs where all aggregator reports are in perfect agreement. As can be seen, for 40% of HYIPs, the maximum allowed investment values are in agreement, while 87% of the time the minimum investment value is reported to be the same by different aggregators.

Of course these attributes are all matters of fact, which the aggregator will have obtained from the HYIP websites (or perhaps from the filling in of a form). However, the aggregators are imperfect and errors are being made. If there was collusion between aggregators and HYIPs then we would have expected to see perfect agreement – either from better channels of communication, or from a consistent set of mistakes being made.

By contrast, when we consider the amount of money that the aggregators report that they have invested into particular HYIPs, we see very little agreement at all:

Aggregator Investment	
Perfect Agreement	0.10
Diversity Index	0.51

Any investment at all allows the aggregators to assess whether the HYIP is paying, and we have just noted there is reasonable agreement about what the minimum value might be. Therefore, we presume that the amounts being invested reflect the initial opinion of the aggregator about the prospects for the HYIP. If there was some kind of universal conspiracy then we would expect to see consistency here, but the aggregators invest the same amount of money into HYIPs in only 10% of cases.

Naturally, even when there isn't unanimous agreement across aggregators, it could still be the case that almost all of aggregators report the same values. Consequently, the two tables also report Simpson's diversity index [1] for each attribute. This measures the similarity of a sample population and it is computed as the sum of the squares of all the probabilities for each attribute value, with a 0 score showing complete diversity and complete uniformity giving a score of 1. Once again, using this measure, we see a high, but imperfect, agreement on matters of fact, but continuing diversity in the investment amount.

3.2 Reporting of HYIP Lifetimes

We now consider the elapsed time between when HYIPs are reported to be created – or at least when the aggregator learns of their existence – and when the HYIPs collapse and the aggregator is no longer prepared to track them.

Figure 2 (left) plots the Cumulative Distribution Function (CDF) of the standard deviations of the reported starting and ending times of HYIPs across aggregators. For around 80% of HYIPs, the standard deviation is very small, at most a few days. However, for the remaining 20% of HYIPs, there is substantial disagreement between aggregators. Furthermore, that disagreement is greater for the last observed time for an HYIP than for the starting time, as indicated by the slightly lower blue dashed line than solid red line in the graph. This is not surprising, given that deciding when to drop an HYIP from results is more of a judgement call for an aggregator than deciding whether to report its existence.

The green dotted line plots the standard deviations for an alternative measure of HYIP collapse. Aggregators keep track of their own investments in HYIPs, reporting each day the cumulative return on investment (ROI). Often, the ROI will ‘flat-line’ – suddenly stop changing – a few days before the aggregator stops tracking the program, because the HYIP is no longer paying out. Hence, we can view the time when the ROI stops changing as an alternative indicator of collapse. As the graph indicates, there is even more variation here – some aggregators stop receiving payments before others. Again, this is not surprising, since HYIPs may not stop all payments at once.

Aggregators generally agree on lifetime, but when there are differences they can be large, so for lifetime value we use the median of the aggregator reports. By using the median (rather than computing the mean), we are better protected against a highly divergent aggregator polluting the overall measure.

Overall, our analysis of aggregator reports is that there is no evidence of collusion, but that their measurements are generally consistent, and that our further analysis based on the median of aggregator values will be robust.

4 The Collapse of HYIP Programs

An HYIP scheme collapses when it can no longer make the interest payments that it has promised. While it may not have completely run out of money, a rational HYIP operator will eventually conclude that paying the next round of interest payments (or refunding someone’s capital) is less lucrative than shutting the scheme down and absconding. These calculations are slightly different in cyberspace than for real world Ponzi schemes because there will be no bankruptcy and no liquidators checking to see if any value can be salvaged from the ruins.

4.1 How Long do HYIPs Survive?

One subtlety in measuring HYIP lifetimes is that some schemes remained viable at the end of our study, making it impossible to observe when these schemes




Fig. 2. CDF of the standard deviations of HYIP lifetimes (left); the graph indicates that aggregators assess similar lifetimes for around 80% of HYIPs. Survival function of HYIP lifetime (right); the graph shows that most HYIPs collapse within a few weeks, but that a small fraction can remain open for several years.

ultimately collapsed. This can be solved using survival analysis: the 187 (12% of the 1576 total) HYIPs that were still operational at the end of our investigation are said to be ‘right-censored’.

A survival function $S(t)$ measures the probability that an HYIP’s lifetime is greater than time t . This is similar to a complementary cumulative distribution function, except that the censored data points must be taken into account and the probabilities estimated. We use the standard Kaplan-Meier estimator [2] to calculate a survival function for HYIP lifetimes.

Figure 2 (right), has a logarithmic x-axis and plots the observed survival function for HYIPs (using the median observed lifetimes across all aggregators). The solid blue line indicates the survival function computed using the HYIP’s last observed time, while the green dotted line plots the survival function using the ROI flat-line method described in the previous section. For very short-lived HYIPs (i.e., less than one week), the lifetime measured using the ROI flat-line method is considerably shorter. However, for longer-lived schemes, the lifetimes are nearly indistinguishable, so we ignore the ROI flat-line method for subsequent analysis, and just use the median of the lifetime values.

The survival function data shows us that while the median lifetime of HYIPs is just 28 days, one in four will last more than three months, and one in ten for more than ten months. That is, although many HYIPs collapse almost immediately, a substantial minority persist for a very long time.

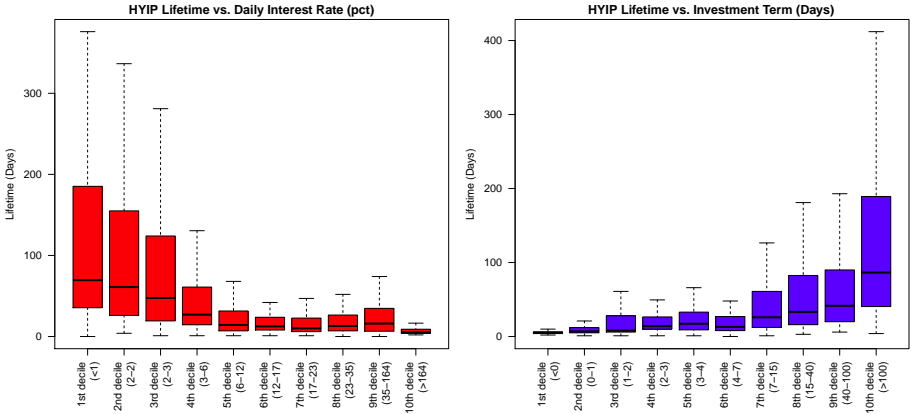


Fig. 3. HYIP lifetimes: HYIPs with lower daily interest rates tend to last longer before collapsing (left); HYIPs with longer mandatory investment periods tend to survive longer (right).

4.2 What Factors Affect HYIP Time-to-Collapse?

Given such regular turnover and wide variation in lifetimes, it is natural to wonder what might prolong or trigger collapse.

Figure 3 examines how the generosity of the HYIP investment terms affects the observed lifetimes. On the left, box plots for HYIP lifetimes are given that span different profit rates. When an HYIP offers a less generous profit rate (less than a few percent daily), there is a greater chance that the HYIP will survive for longer. Once the profit rates become more outlandish (such as the half of HYIPs offering more than 10% daily returns), HYIP lifetimes are more consistently short. We conclude that the offering of higher rates of return does not bring in sufficient investment to offset the cost of servicing existing commitments.

Another factor is the minimum investment period required by the HYIP. Figure 3 (right) plots HYIP lifetimes sorted by investment term. As expected, HYIPs that require longer investment terms tend to be more stable. However, we note that there is still substantial variation in lifetime even for less generous interest rates and longer investment terms. Evidently, some HYIPs cannot attract enough investment to sustain even these more modest programs.

4.3 Can Users or Aggregators Predict Collapse?

Several aggregators rate HYIP ‘quality’, often on a scale of zero to five stars. The rating can vary considerably over time, ostensibly according to the risk level associated with the scheme. Some aggregators also compile user ratings, typically collected in the form of positive (and sometimes negative) votes. We now examine how the crowd’s rating compares to that of the curator’s.

We focus on the four aggregators that report both user and aggregator ratings on a finite scale. Some aggregators simply tally the total number of user votes, while others report the absolute difference between positive and negative votes. We exclude these reports from our analysis to ease comparisons. The ratings we study are based on a score of zero to five, zero to ten, or out of 100; we normalize all ratings to a percentage.

While the ratings have been collected throughout an HYIP's lifespan, we have decided to take a closer look at the rating given 7 days prior to each HYIP's collapse. A low rating issued at this point would indicate to prospective investors that the bottom will soon fall out (if it has not already). The results are given in the following table.

Aggregator	# HYIPs	User Rating			Aggregator Rating		
		Avg.	$\leq 50\%$	$\geq 80\%$	Avg.	$\leq 50\%$	$\geq 80\%$
everyhyip.com	46	87%	13%	85%	20%	83%	9%
hyip.com	265	52%	45%	47%	8%	94%	3%
hyipranks.com	107	96%	4%	96%	35%	89%	7%
hothyips.com	292	50%	48%	46%	32%	92%	0.3%
Average	-	60%	38%	57%	22%	92%	3%

Overall, user ratings are consistently much more positive than aggregator ratings. Consider the ratings for `hyipranks.com`: the average user rating one week prior to collapse is 96%. Across all HYIPs, 96% were awarded a user score of 80% or higher, but only 4% had a score below 50%. By contrast, the average assessment directly issued by `hyipranks.com` is only 35%. Moreover, 89% of HYIPs are given a low score, compared to just 7% that receive a score over 80%.

Why do we see such divergence in ratings? Those who have already invested in an HYIP have a very strong incentive to attract new investors. Consequently, they are highly motivated to vote early and often in support of their investment. The aggregators, on the other hand, fully expect HYIPs to collapse and must provide more accurate assessments in order to gain the trust of visitors. Viewed in this way, it is not surprising that the crowd will not accurately predict collapse.

5 The Role of Digital Currencies

Digital currencies are an essential component of a functioning HYIP ecosystem. They allow investors to convert local hard currency into a multi-national form that is suitable for transfers to and from the HYIP. Occasionally an HYIP will directly accept wire transfers or credit card payments. However, this is unusual because if the HYIP operator works within the traditional financial system, then they risk being identified when the fraud collapses, and they will be less sure that they will be able to hang on to any profits.

We found that 22 currencies were accepted for use at the HYIPs we tracked. Most of these were only offered by a handful of HYIPs (including 14 HYIPs that took PayPal, 7 Moneybookers and 1 Western Union). We list the six most

common currencies below and note that the most common, by far, were Liberty Reserve and Perfect Money, accepted by 83% and 70% of HYIPs, respectively. Both currencies are based in Central America.

Currency	HYIPs		Country	% HYIP Backlinks
	#	%		
Liberty Reserve	1 309	83%	Costa Rica	33%
Perfect Money	1 095	70%	Panama	72%
AlertPay	397	25%	Canada	10%
SolidTrustPay	51	3.2%	Canada	60%
Pecunix	21	1.3%	Panama	81%
GlobalDigitalPay	20	1.3%	Hong Kong	71%

Digital currencies are riskier than traditional currencies for both the investor and the HYIP operator. When the time comes to cash in and convert back to hard currency, the exchange rate may have changed significantly, or there may be no liquidity – if many of the customers of a digital currency simultaneously ask to cash in their holdings, then the currency’s operators may not be able to pay up (e.g., the HYIP-associated StrictPay currency appears to have collapsed in this way [3]).

The digital currencies that HYIPs accept have terms and conditions that forbid their use with HYIPs. This creates the additional risk that assets could be frozen or confiscated for violating the rules. Furthermore, any digital currency that facilitates widespread criminality runs the risk of being shut down by law enforcement, as happened to e-gold [4].

Liberty Reserve has a warning on its website advising against investing in HYIPs, noting that payments are ‘non-revocable’ and that they cannot be held liable for fraudulent activities by its users. Such admonishments raise the question: how much of these currencies’ profits come from HYIP activity?

We attempt to shed light on this by examining the backlinks from other websites into the currency websites. We used Yahoo Site Explorer³ to gather 1 000 backlinks for each of the most common currencies and calculated what proportion of the incoming links came from HYIP-related websites. The results are listed in the right-most column of the table.

72% of the backlinks to Perfect Money are from HYIP-related websites, as are 33% of the backlinks to LibertyReserve. This leads us to conclude that a substantial proportion of the revenue to these currencies comes from HYIPs. Note that for AlertPay, the third-most popular currency, only 10% of the incoming links are from HYIPs. Indeed, many of AlertPay’s other incoming links are from legitimate businesses, such as the web-hosting company `proddhosting.net`, which uses AlertPay to process payments. AlertPay is based in Canada, and that may mean that they are more easily pressured by first world regulators, than the currencies based in Panama and Costa Rica.

³ <http://siteexplorer.search.yahoo.com>

6 Policy Options for Disrupting the HYIP Ecosystem

One of the first questions to ask when considering policy interventions into online scams is how prevalent the scam is. If only a few people are affected, then the criminality may not be worth pursuing, especially when – as in this case – many of the investors are aware that the sites are fundamentally fraudulent.

It is difficult to directly measure how many people and how much money are invested in HYIPs. However, we can use some publicly available proxies to derive an order-of-magnitude estimate of HYIP impact.

As part of its Adwords program, Google offers a Keyword Tool that returns similar search phrases to those given as input.⁴ We entered the phrases “hyip” and “high yield investment program”, and were returned 100 closely related phrases. Google also offers a related service called Traffic Estimator that estimates for any phrase the number of global monthly searches. We plugged all 102 HYIP-related phrases into the tool to arrive at an estimate of 441 000 monthly searches for these terms on Google.

We can use this value to create a rough estimate of the monthly investment levels to HYIPs using the following formula:

$$\frac{\$ \text{ HYIP invest}}{\text{month}} = \frac{\# \text{ Google mo. searches}}{\text{Google market share}} \times \% \text{ invest} \times \text{invest amount}$$

Google’s global market share in search is known to be 64.4% but the other terms in this equation are much harder to estimate. We do not have reliable data on the fraction of users who learn about HYIPs that ultimately invest, or how much money they put in. A plausible, conservative, guess is that at least 1% of people who search for HYIPs go on to invest in an HYIP.

Researchers investigating spam-advertised pharmaceuticals found that 0.5% of site visitors added items to their shopping carts [5], while in an earlier study they found an approximately 8% conversion rate for non-pharmaceutical goods [6]. Leontiadis et al. estimated that between 0.3% and 3% of people looking for drugs via web search ultimately purchased the goods from illicit retailers [7]. While the investment rate for HYIPs could undoubtedly differ from that for pharmaceuticals, these data points do suggest that a 1% conversion rate for HYIPs is plausible.

From observation of the statistical information that some sites provide, we will guess that the average investment is \$1 000. Plugging these numbers into the above formula we estimate that HYIPs attract at least \$6 million per month in revenue.

Given that around 600 000 people search for HYIPs each month, we conclude that HYIPs are indeed a substantial scam worthy of policymakers’ attention. So what should be done? We now consider a range of interventions and assess their likely impact.

⁴ <https://adwords.google.com/select/TrafficEstimatorSandbox>

Option 1: Engage Law Enforcement. Given that HYIPs are illegal in nearly all jurisdictions, it is logical to seek the support of law enforcement. In the US, the Commodity Futures Exchange Commission (CFTC) has been given the power to enforce violations of the Commodities Exchange Act of 1936. The CFTC regularly uncovers Ponzi schemes whose perpetrators and victims are based in the US. International cooperation is possible: the CFTC recently arrested Jeffrey Lowrance and extradited him from Peru for allegedly running a Ponzi scheme that solicited investors via the Internet [8]. Consequently, engaging the CFTC could lead to successful criminal prosecution.

However, the usual warnings about prosecuting online crime [9] apply: collecting evidence across international borders is difficult, slow and expensive; the perpetrators may be located in countries unwilling to cooperate. Google's data shows that that 85% of HYIP-related searches are made from outside the US, so victims will be spread across the globe, necessitating an international response.

Policy interventions that apply pressure to key intermediaries have historically been one of the most successful ways to address illicit activity online. For example, the US Unlawful Internet Enforcement Act of 2006 has largely eliminated online gambling by US residents by requiring payment processors to block credit-card payments to offshore gambling sites. The Digital Millennium Copyright Act of 1998 created a notice-and-takedown regime whereby online service providers receive immunity for complying with take-down requests issued by copyright holders. So we now turn to considering potential intermediaries that might be enlisted to disrupt the HYIP ecosystem.

Option 2: Target Digital Currencies. The digital currencies that HYIPs rely on for customer accounts are a logical target. As shown in Section 5, a handful of currencies facilitate most HYIP transactions. The biggest offenders (Liberty Reserve and Perfect Money) are undoubtedly aware of their role in funding HYIPs, so bringing it to their attention is unlikely to make any difference. The banking regulators in their claimed home countries (Costa Rica and Panama) might be persuaded to cooperate with an outside crackdown. However, even if they stopped processing HYIP payments, it is likely that alternative currencies would come to the fore.

Option 3: Squeeze Credit-Card Payments. Another option is to block the funding of digital currencies by credit cards. At present, a credit card can be used to fund the most popular digital currencies, including Liberty Reserve and Perfect Money. Although this is an obvious opportunity to apply pressure, it might prove difficult to identify all the intermediaries that can supply the currency, and ultimately the traffic would shift to wire transfers instead.

Option 4: Undermine Aggregators. A more promising approach is to disrupt the aggregators, since they are essential for establishing trust in HYIP transactions.

For example, one could target the registrars that have registered the domains in use. Persistent websites are essential for establishing the reputation of the aggregators, so they are more likely to be adversely affected by a domain name

seizure than, say, malware-distributing sites. Many aggregators are currently served by North American companies (e.g., `hyip.com` and `hyipranks.com` are registered through GoDaddy, `maxhyip.com` is on Tucows, and `hyip.com` is for sale by American domain-parking firm Sedo). However, this is likely to require new legislation, since the aggregators are merely describing and linking to the HYIP sites. It could be some time before such legislation was in place in the USA, let alone in all the jurisdictions to which the sites could move.

Alternatively, the aggregators' income stream could be disrupted. Four of the nine aggregators we studied – `hyip.com`, `iehyip.com`, `hyipranks.com` and `hyipinvestment.com` – are members of the Google Display Network. Google, and the other advertising networks, might choose not to work with sites that knowingly link to fraudulent sites. Whether stopping this source of income would cause all the sites to close cannot be known for certain, but it is relatively straightforward, and arguably in the best interests of the advertising networks to cease their financial association with criminality comparison sites.

Option 5: Target most the successful HYIPs. A final option is to attempt to expedite the demise of HYIP websites. While this might appear a hopelessly difficult task given that we have observed around 1 600 HYIPs in just nine months of data collection, targeting the small number of long-lived HYIPs could be effective. A long-lived HYIP is bound to be continuing to attract many victims, since new recruits are needed to prolong the life of the scam. Consequently, efforts to disrupt these programs are very likely to reap substantial rewards.

Over one third (49) of the 141 HYIPs that have been online for more than six months are registered via eNom, a US-based registrar. 26 are registered through Indian-based Directi, along with another 14 on US-based GoDaddy. Consequently, making registrars aware of the criminal behavior being facilitated by these websites could trigger a short-term disruption.

On balance, while each of the discussed options may help, we expect that options 4 and 5 are likely to be most helpful for disrupting the current HYIP ecosystem. We also believe that action by law enforcement (option 1) could do a lot of good in the longer term.

7 Related Work

During the past decade, online criminality has proliferated [9]. In response, a number of measurement studies have quantified various frauds and recommended suitable interventions. Of particular relevance are studies that examine user susceptibility to various scams, such as fake antivirus [10,11] and extortionate social-engineering scams [12]. Stajano and Wilson identify seven principles common to offline scams that often translate into online scams [13]. At least five of these principles apply to HYIP investment: the herd principle (false safety in numbers), the dishonesty principle (victim's own illegal behavior held against him), deception principle (things are not what they seem), need and greed principle (desperation increases vulnerability), and the time principle (time pressures

increase bad choices). Consequently, while we believe that many HYIP investors are likely to be aware of the fraudulent nature of their investment, they are nonetheless being masterfully deceived by con artists. Furthermore, it is entirely plausible that some victims fully believe in the legitimacy of their investment.

The use and abuse of digital currencies has been examined since the inception of the Financial Cryptography conference. Optimists pointed to the potential to enhance revenue [14] or freedom through anonymity [15]. However, even in these early days, others fretted about the potential for abuse of digital cash, such as money laundering [16]. More recently, Anderson identified non-revocability as the key feature of digital payments that appeals most to online criminals [17]. Indeed, the non-revocability of payments issued in the currencies underpinning the HYIP ecosystem is essential for its successful operation.

The security and reliability of crowdsourcing in information security applications has been investigated by Moore and Clayton [18] (phishing) and Chia and Knapskog [19] (web security). These papers discuss the distinct challenges of crowdsourcing applications when participants may be motivated to lie, as we have found for users promoting flagging HYIPs.

A final area of relevant work is in the examination of interventions to combat online crime. In an expansive study of goods advertised by email spam, Levchenko et al. [20] found substantial concentration in the registrars used by spam-advertised websites. They also found that only 3 banks processed the bulk of payments. We report similar levels of concentration in the HYIP ecosystem. Clayton described how shutting down hosting providers that facilitate spam transmission can have a disruptive short term effect [21]. Finally, Liu et al. [22] examine the prospects of enlisting registrars to suspend ‘known bad’ domains, concluding that the criminals are more adept at shifting to new domains faster than the offending domains can be suspended. While this may be true for domains used in email spam, we are more optimistic for registrar-level intervention in combating HYIPs due to the persistence of successful schemes.

8 Conclusions and Future Work

We have presented the first detailed analysis of HYIPs – fraudulent online Ponzi schemes. We have provided some baseline measurements by leveraging data from the aggregator sites that exist to help investors pick where to place their money. We have shown that the aggregators are basically truthful, and used their data to show that HYIPs last longer with lower interest rates delays before payments are made. We have also shown that the aggregators are better than ‘the crowd’ in warning of HYIP collapse, which we believe is directly related to the crowd actively wishing to hype the prospects of the HYIP they are invested in.

Nonetheless, this paper has only scraped the surface in measuring and understanding HYIPs, and there is much more data to collect and process. It is already clear to us that many of the sites are related to each other as criminals create new instances to replace HYIPs that have collapsed. We have been un-

able, so far, to use WHOIS data to identify serial offenders but we expect to make headway when we consider the structure and content of the websites.

We are particularly interested in the subset of HYIPs that provide a running commentary on the number of accounts opened, and the sums of money being invested, withdrawn and paid out as interest. We hope to use these to build a better model of HYIP collapse, and provide better estimates of the sums of money passing through these criminal enterprises.

As we extend our analysis and measurement of harm, we intend to ensure that this paper's other key contribution – a detailed analysis of how this criminality might be disrupted – may be of even greater relevance to policy makers.

References

1. Simpson, E.H.: Measurement of diversity. *Nature* **163** (1949) 688
2. Kaplan, E., Meier, P.: Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association* **53** (1958) 457–481
3. Lorenzini, M.: Strictpay scam: Thoughts before strictpay shutdown. HYIP News (June 2010) <http://www.hyipnews.com/news/17190/STRICTPAY-SCAM-THOUGHTS-BEFORE-STRICTPAY-SHUTDOWN/>.
4. Zetter, K.: Bullion and bandits: The improbable rise and fall of e-gold. *Wired* (June 2009) <http://www.wired.com/threatlevel/2009/06/e-gold/>.
5. Kanich, C., Weaver, N., McCoy, D., Halvorson, T., Kreibich, C., Levchenko, K., Paxson, V., Voelker, G.M., Savage, S.: Show me the money: Characterizing spam-advertised revenue. In: *Proceedings of USENIX Security 2011*, San Francisco, CA (August 2011)
6. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G., Paxson, V., Savage, S.: Spamalytics: An empirical analysis of spam marketing conversion. In: *Conference on Computer and Communications Security (CCS)*, Alexandria, VA (October 2008)
7. Leontiadis, N., Moore, T., Christin, N.: Measuring and analyzing search-redirection attacks in the illicit online prescription drug trade. In: *Proceedings of USENIX Security 2011*, San Francisco, CA (August 2011)
8. Commission, C.F.E.: Press Release PR6074-11: CFTC charges Jeffery A. Lowrance and his company, First Capital Savings and Loan, with operating a million dollar foreign currency Ponzi scheme (July 2011) <http://www.cftc.gov/PressRoom/PressReleases/pr6074-11.html>.
9. Moore, T., Clayton, R., Anderson, R.: The economics of online crime. *Journal of Economic Perspectives* **23**(3) (Summer 2009) 3–20
10. Cova, M., Leita, C., Thonnard, O., Keromytis, A., Dacier, M.: An analysis of rogue AV campaigns. In: *Proc. RAID 2010*, Ottawa, ON, Canada (September 2010)
11. Stone-Gross, B., Abman, R., Kemmerer, R.A., Kruegel, C., Steigerwald, D.G., Vigna, G.: The underground economy of fake antivirus software. In: *10th Workshop on the Economics of Information Security*, Fairfax, VA (June 2011)
12. Christin, N., Yanagihara, S., Kamataki, K.: Dissecting one click frauds. In: *ACM Conference on Computer and Communications Security (CCS)*, Chicago, IL (October 2010) 15–26
13. Stajano, F., Wilson, P.: Understanding scam victims: seven principles for systems security. *Commun. ACM* **54** (March 2011) 70–75

14. Birch, D.G.W., McEvoy, N.A.: Electronic cash - technology will denationalise money. In Hirschfeld, R., ed.: *Financial Cryptography*. Volume 1318 of *Lecture Notes in Computer Science.*, Springer (1997) 95–108
15. Chaum, D.: Achieving electronic privacy. *Scientific American* (August 1992) 96–101
16. Wayner, P.C.: Money laundering: Past, present and future. In Hirschfeld, R., ed.: *Financial Cryptography*. Volume 1318 of *Lecture Notes in Computer Science.*, Springer (1997) 301–306
17. Anderson, R.: Closing the phishing hole: Fraud, risk and nonbanks. In: *Federal Reserve Bank of Kansas City – Payment System Research Conferences*. (2007)
18. Moore, T., Clayton, R.: Evaluating the wisdom of crowds in assessing phishing websites. In Tsudik, G., ed.: *Financial Cryptography and Data Security*. Volume 5143 of *Lecture Notes in Computer Science.*, Springer (2008) 16–30
19. Chia, P.H., Knapskog, S.J.: Re-evaluating the wisdom of crowds in assessing web security. In: *Financial Cryptography and Data Security*. (2011)
20. Levchenko, K., Chachra, N., Enright, B., Felegyhazi, M., Grier, C., Halvorson, T., Kanich, C., Kreibich, C., Liu, H., McCoy, D., Pitsillidis, A., Weaver, N., Paxson, V., Voelker, G., Savage, S.: Click trajectories: End-to-end analysis of the spam value chain. In: *IEEE Symposium on Security and Privacy*, Oakland, CA (May 2011) 431–446
21. Clayton, R.: How much did shutting down McColo help? In: *Sixth Conference on Email and Antispam (CEAS)*. (July 2009)
22. Liu, H., Levchenko, K., Felegyhazi, M., Kreibich, C., Maier, G., Voelker, G.M., Savage, S.: On the effects of registrar-level intervention. In: *USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, Boston, MA (March 2011)