Crime Markets

Does Truthcoin condemn us to a lawless plutocracy?

Paul Sztorc
truthcoin@gmail.com
Version 1.0 - November 30th, 2013[1]

Summary

Some worry that censorship-resistant prediction markets will be used to encourage assassinations (and other crimes); this concern does not hold up to a sober examination. "Assassination markets" (AMs), as originally proposed by Jim Bell, are irreconcilably different from Prediction Markets (PMs). My experimental method for funding public goods with PMs has features which render it incompatible with crime. Furthermore, markets would generally present an excessively-complex, risky, and convoluted form of criminal financing. Truthcoin presents a (peaceful) alternative for accomplishing ideological goals, which features greater persuasiveness as well as lower cost. I conclude with a short discussion covering [1] recourse for those affected by AMs (of any kind), [2] features of Truthcoin designed to amplify the inherent impracticalities of AMs, and [3] the (necessarily relevant) *total* net effect of Truthcoin on political assassinations, general crime and general human welfare.

[1] I claim ethical kudos for writing this essay *before* publishing the Truthcoin Whitepaper. In my view, it is very revealing that other projects (and alphabetic neighbors), argue for 'general censorship-resistance' without doing likewise.

The Problem

Anonymous Crime Without Recourse

Protocols are immortal. Once they are on, we can't shut them off[2]. If we want a protocol to go away, we have to just hope that no one finds it useful.

Truthcoin is a protocol (and therefore immortal). Truthcoin is designed to resist censorship (the mechanism will try to answer almost *any* question put to it). Thirdly, I earlier described a method for using Truthcoin to finance certain goods and services.

Combined, these features yield an obvious concern: "What if Truthcoin is used to finance something that I don't like? I can't censor Truthcoin, or shut it off!" This concern is sometimes expressed specifically as "assassination markets", where (for example) the canonical PM "Will Person A be elected president in 2016?" is perverted into "Will Person A be killed in 2016?". Then, supposedly, violent murderers can "bet" on [No] from the comfort of their own homes. When a sufficient total sum of money has been bet on [No], people will begin to notice that, if Person A happened to somehow die, anyone who had bet on [Yes] stood to make a lot of money. The theoretically inevitable conclusion is that the small donations from many individuals will produce a situation where (at least) one person is willing-and-able to pull the trigger, while individually placing a large single bet on [Yes]. The general inference is that prediction markets enable a kind of "crime store" where ordinary people can buy assassinations, or other horrible things, such as bribes or acts of terrorism. What's more, the assassins themselves have some built in moral and legal defenses ("I was unfairly tempted by the high quantity of money offered to me." "If I didn't take advantage of this opportunity, someone else would have.") Society is therefore harmed, irreparably, by this tide of new violence.

[2] It may be more accurate to say "once they are learned, they can't be forgotten", as protocols are merely a set of rules. While a set of rules may be yet-undiscovered, it has, in a sense, always "existed" and will always exist.

Problems with the Problem

I intend to show that the seemingly-plausible story presented above does not hold up to a sober examination. First, I'll warm up by introducing the original design for "assassination markets", and show that prediction markets and assassination markets are irreconcilably different. Second, I review the previously-introduced concept of a trustless dominant assurance contract (T-DAC), and elaborate on how this contract can be used to encourage real-world behaviors, *only if* it is given appropriate Schelling Indicators (SIs). Thirdly, I will argue that the SI, in practice, will need to be publically set *after* the trade has been made. This necessarily either [1] leaks the identity of the agent manipulating the SI or [2] gives an hour's advance warning of the imminent crime. Fourth, I will list ways in which a criminal T-DAC ("Crime Market"), even if it could exist, would be inferior to existing methods of crime. Fifth, I demonstrate that the use of Crime Markets to effect political change (the stated purpose of AMs) is largely missing the point: Prediction Markets themselves can already (peacefully) accomplish these ideological goals, not only with greater persuasiveness, but also with the opportunity for activists to actually *earn* money (as opposed to *spend* money in a "Crime Store"). Sixth, I lay out an extreme (but intimidating) counterstrategy for those "targeted" by such an attempt. Seventh, I describe how I have edited the Truthcoin proposal to amplify all of the previously-discussed disadvantages. I close by mentioning that, even if one ignores the entire analysis, and just assumes that (somehow) Truthcoin results in more assassinations, we can be reasonably certain that it will save so many other lives that, net, Truthcoin's existence would still be an ethical imperative.

1. Prediction Markets are not Assassination Markets

The phrase "assassination markets" has a very specific meaning. It refers to an unsavory[3] essay entitled "Assassination Politics" by Jim Bell which hypothesizes a smart contract[4] designed specifically to facilitate assassinations.

First, let me standardize the terminology: for the remainder of this essay I will refer to the doomed individual as the "Victim", those who attempt to purchase the assassination as "Contributors"[5], and the agent who carries out the execution as the "Assassin".

Bell's Assassination Market (AM) proposal is for [1] the selection of a Victim and creation of the relevant smart contract, [2] the collection of monetary contributions from the Contributors, thus funding the contract, at which point [3] an Assassin decides to fulfill the terms of the contract (and kill the Victim). He does this by [4a] entering a "prediction" on when the individual will die, and then [4b] killing him/her on that date. If all goes as planned, [5] the assassin gets away with the crime and collects a large payout of anonymous e-cash (as the Victim just happened to die on the exact date the Assassin predicted).

A Prediction Market (PM) is supposedly a version of this: if the topic of the PM is "Will Victim die before date D?" then Contributors can purchase shares of 'No', and, if enough money is on the line, the Assassin can purchase a corresponding about of 'Yes' and then kill the Victim, becoming wealthy when his prediction turns out to be accurate.

However, there is an important difference. With the AM, the Assassin had the *unique* ability to choose (and keep secret) the date of the assassination in advance, and so the Assassin could

[3] Bell's evocative characterization of his service as "universal Second Amendment rights" notwithstanding, the attempt to square [1] the use of this service against (innocent) civil servants against [2] a "pre-emptive" non-aggression principle is entirely ridiculous (primarily because, unlike in an alley mugging, the citizen and civil servant would know so little reliable information about each other).

[4] In your author's opinion, certain other projects are *far* more amenable to AMs, and yet these projects mysteriously escaped similar criticism (until very recently).

[5] The more-appropriate term "Murderers" might lead the reader to confuse these individuals with the Assassin (the agent pulling the trigger).

therefore virtually guarantee that he and only he would collect the funds. With the PM, anyone can purchase any amount of 'Yes' at any time.

When "dies *on a specific date*" blurs outward into the overwhelmingly general "dies at some point", the similarities between AMs and PMs break down completely. Consider these phenomena: [1] a rival Assassin spies on the actual Assassin and steal his payout by front-running his trade, [2] the Vitcim bets on his own death (removing the incentive) and even pretends to have been killed (profiting from his enemy's contributions). Or, [3] Any informed third-parties (co-conspirators of the Assassin, or friends of the Victim) might wait for the Assassin's pro-death 'Yes' trade to go through, themselves make a (now-equally-profitable) pro-life 'No' trade, and then foil the Assassin's plan (by ratting him out to the authorities in exchange for legal immunity, or using personal resources to protect the Victim). These counterexamples all exploit the fact that *anyone*, regardless of intent, can trade on either side of the single degree of freedom[6].

In summary, Bell's "Assassination Markets" aren't really "markets" in the financial sense: there's no price, no trading activity, and no speculators or info-aggregation. In fact, you can't even *sell* in this market. AMs are only "markets" in the 'store' sense.

Can someone modify PMs to make them more like AMs?

2. The Circularity Problem -- Public Markets vs. Private Incentives

The use of a two-state PM as an AM creates a paradox. As Contributors place more money on 'Live': [1] they *encourage* someone to bet-on-and-cause 'Die', [2] they imply that 'Die' is becoming *less likely*, and yet, [1+2] any *encouragement*, by definition, causes something to become *more likely*.

[6] A market with two states has one "degree of freedom" because the price of State 2 is always "1 – the price of State 1". As such there is only one signal, and "inducements" to bring about the event cannot be separated from "forecasts" about the likelihood of the event.

Markets are public institutions, not private ones, and, unlike in a store or a casino, the odds (prices) are *not fixed* but *react strongly* to the wagers as they are placed.

“Now that the unemployment situation has improved, stock prices have nowhere to go but up.”

–Someone who knows nothing about Finance

“Now that there’s 50 million dollars on ‘Live’, the likelihood of ‘Die’ has nowhere to go but up.”

–Those concerned about Prediction-to-Assassination Markets

AMs (not PMs) avoid this distinctive market-reactivity by outright rejecting the market format for something else (in which the monetary reward is privately collectable). Similarly, in the T-DAC (above) I make two modifications to the PM format to try and achieve something similar: [1] banning sales (one can only buy, in this market, and later redeem shares –a very minor limitation– (as one can pseudo-sell by purchasing until they own all outcomes equally), and [2] introducing the more general/arbitrary “Schelling Indicator” (SI). This change is the additional degree of freedom required to prevent *encouragements* of an event from reducing that event’s *likelihood*.

“How much can the Assassin collect?”

Let’s talk about why is such a crucial requirement. In a two-state Truthcoin-PM (without SIs), the maximum possible amount of money that can be extracted from a Market at a given time is given by $B(t) = \max\big(q_1(t) - q_2(t)\ ,\ q_2(t) - q_1(t)\big)$ where $q_x(t)$ represents the total outstanding shares of state x at time . This assumes that, somehow, these shares can be acquired for free (ie purchased at a price of zero).

For example, if, at t=3, there were 200 outstanding shares of “live”, and 0 of “die”, then an Assassin could make at most $B(3) = \$200$ by buying 200 “Die” shares (for $0), causing the outcome to be “Die”, and then redeeming each share for $1. If, at t=60, there were 7,500,000 shares outstanding of “live” and 8,400,000 shares outstanding of “die”, then at this point in time the

Assassin can earn at most $B(60) = \$900{,}000$ by buying 900,000 shares of "Live" and forcing the outcome to be "Live". Some notation, for brevity: $B(3) = \$200_{Die}$ and $B(60) = \$900{,}000_{Live}$.

If the smart contract in question were *only* a smart contract, and not also a Prediction Market (where *encouragements* affect *likelihoods*), things would remain in this state of affairs.

The Centrifugal Governor

One would hope that any critical reader would be immediately suspicious; after all, we live in a world that has all kinds of markets, yet nothing is ever purchased in such a roundabout way. There is a reason: whenever $B(t)$ is increased, $B(t)$ is immediately induced to *decrease*.

Why is this the case? For starters, it is easy to increase $B(t)$, by spending additional money on "Live", at t=3, to increase $B(3)$ from $\$200_{Die}$ to, perhaps, $\$40{,}000{,}000_{Die}$ (at cost $(\$40{,}000{,}000 - \$200) * p_{Die} = \$39{,}999{,}800$). In this instant, $B(3) = \ \$40{,}000{,}000_{Die}$ might appear to be overwhelmingly threatening…after all, anyone who can cause this individual's death can bet on "Die" and collect up to 40 million dollars!

However, that ominous incentive is dwarfed by a new one which has just emerged. The price of "Die", thanks to the volume of money on "Live", has now asymptotically approached 0%. Such a low price would seem to indicate that the assassination has a very low likelihood (recall that in PMs, likelihoods are prices). Yet, with a $40 million incentive to perform the assassination, it would paradoxically be very likely. The market now contradicts itself.

The market will quickly self-correct, and resolve this contradiction. Notice that **roughly 40 million shares of "Die" are now *nearly all* available at the price of zero**. While the Contributors paid a lofty $40 million to establish the first set of shares, **the countervailing set of shares are available almost for free**.

"Of course", you may say, "this is so that the Assassin may claim his reward." And so it would be if *only the Assassin* could make this trade (as in Bell's AMs). Yet (in PMs) **anyone can make this trade.**

"Who might make this trade?" Examine the costs and benefits in greater detail: $B(3) = \$40{,}000{,}000_{Die}$, so there are 40 million shares of "Die" which can be purchased for nearly zero (again, this is because Markets, unlike Casinos or Retail Establishments, have prices which *react* to the wagers placed). The cost is therefore quite small, and widely affordable. And the benefit is a lofty 40 million dollars, should the Victim die for any reason (car crash, heart attack, drug use, etc.). The life insurance agencies of the world (or anyone with access to actuarial tables) would jump at such an opportunity. Even if we allow a market with multiple equilibria, we can know with confidence that $\boldsymbol{B(3)}$ **will be kept away from high values automatically (for free)**, and converge to some safe, boring, lower value.

Behold the fundamental contradiction: it is *impossible* for PMs to both [1] accurately measure the likelihood of something, and [2] change the current likelihood from the current market prices to something else. As the incentive to murder increases, mathematically the likelihood must decrease: the PM environment, by definition, shackles the two concepts to opposite ends of the same chain. Yet, *by definition* an incentive makes something *more* likely, not less. So now the two ends of the chain are shackled together in a big circle…pulling on the chain merely rotates it, making no net progress whatsoever.

Market Microstructure Games

Prove-ably Betting on Everything

What if Contributors do not buy shares of "Live", and instead only purchase increases in liquidity parameter *b*? This is a commitment to thicken the entire market (a uniform offer to purchase all shares at all prices).

Now, exactly as planned, anyone who can alter the likelihood of death (in any direction) has a thick market from which to extract revenues. It is no longer cheap to drain the charged “battery”, which can in fact accumulate an unlimited quantity of funds.

Of course, while this does allow the Market to store up and accumulate funds, it does nothing to change the circularity problem: if it really is the case that a more liquid Market creates an assassination incentive, then third parties should force the odds (price) of “Die” down, counteracting the original incentive (and warning the Victim that it is time to consider something like witness protection). It is still the case that a lot of money on “Live” encourages third parties to take a chance on “Die”, as “Die” is so much cheaper, (though not nearly-free) and offers such a disproportionate reward. So $B(T)_{Die}$ is still kept away from extreme values.

Secondly, while “Die” shares are no longer free for anti-Assassins to pick up, they are now correspondingly inconvenient for the actual Assassin to acquire. The Assassin will need to own (and risk) a great deal of working capital in order to get paid. Assassins are unlikely to be wealthy (wealthy people have much more, absolutely and relatively, to lose by going to prison), but, in the alternative case (where the money is supplied by another), one wonders why, if the Assassin plans to do something illegal anyway, he doesn’t simply backstab his criminal-lender and steal the loaned working capital for himself. Likely a much easier way to make money, and one can’t be sued in court by someone with unclean hands. As trust has been reintroduced, one wonders, “Why bother with all of this cryptography in the first place?”.

A third (severe) disadvantage to contributing via liquidity-increase is that, whereas previously the price of “Live” would always be near 1 (such that Contributors would get their money back if the Vitcim didn’t die, thus fulfilling the requirement of an assurance contract), in this case the money spent by the Contributor is lost forever. It is no longer an assurance contract, nor even a “Crime Store” (in stores you always get what you pay for), instead it is something more akin to a claw machine. Any Victim, can acquire this death-bounty (without having to fake his death) by buying "Live" as the Market is fed liquidity, and successfully *living*. The Victim is then, in effect, paid a sum of money to improve his health and take care of his safety. Such an outcome is entirely opposite the one which Contributors originally intended (to offer a payment to someone who could *reduce* the Victim's health).

Zooming Out, to Zoom Back In

Let's consider something else. What if a *group* of people spends 10 million dollars, buying 10 million pairs of {one "Live" share + one "Die" share}, and then they collectively offer to *individually* swap the set of "Die" shares[7] in exchange for 5 million dollars. Now, it is (unfortunately) costly for anti-Assassins to purchase the (eventually worthless) "Die" shares, but an Assassin would find the "Die" shares to be an opportunity to double their money (provided that this Assassin has an extra 5 million dollars lying around, more on this later).

First of all, if this did encourage assassinations (for any reason), it would *still* run into the circularity problem. The public price of "Die" would be pushed up, past 50%, encouraging some of the people owning "Die" to cash out these shares for immediate profit. Non-Assassin speculators would pick up or hold onto "Die" shares, hoping to double-up without doing any criminal work. Even stranger, as some members of this group transferred "Die" shares to the Assassin, other members would have a net incentive to keep this individual alive (so that their "Live" shares pay off). The battery, once charged, always starts draining itself.

To move further, let's consider the case where some *individual* spends $10 million on 10 million shares each of 'Live' and 'Die'. He then offers up the 10 million 'Die' shares for 50 cents each, and refuses to transact any but the *whole total* to anyone except *one person*. The Assassin can double his money, and so he buys, but other people cannot double their money, so they don't buy.

[7] I'm not sure that Truthcoin will even have this feature.

Eventually the Assassin ends up being the sole owner of a huge quantity of shares, and pulls the trigger.

While this does, finally, get the job done, the *original purpose* of AMs was to [1] allow Contributors to band together as a group and pool their economic resources, and [2] allow the (hitherto unknown) Assassin to opt-in to the contract. If we reduce Contributors and Assassins from "a large vague team that includes everyone" to "two specific and known individuals", then the crypto-infrastructure becomes superfluous: one individual could just pay the other the 10 million directly (as with any other job –criminal or otherwise). In other words, the complicated Crime Market infrastructure is vastly inferior to existing escrow technologies (multisignature, private-bank, criminal-lawyer, etc). Why would the Contributor invoke the blockchain, and it's permanent, public paper trail? Why should the Assassin have to come up with ~5 million dollars of working capital in order to even qualify for this job? Why place this money in the hands of the decentralized Oracle? By doing these things, needless complexity has been introduced to a system which itself lacks each of the original "Assassination Market" features.

Another peculiarity is that, if one is going to use a convoluted PM to purchase crimes, a great deal of critically-important criminal information is going to be made public: namely, all of the details of the crime. On top of that, given the importance of this large payment, the current technology requires a 1 hour settlement period. In other words, **the police will know about this exact crime at least an hour before it takes place**. Fundamentally, the trade can't be hidden, because the Victim himself can query the buy opportunity (to check and see if it is still available).

Real World Parallels

Market Efficiency Trumps Everything

There's no denying it: Markets reward those who predict the future, and so they will necessarily also-reward individuals who *create* the future.

Do market-prices *more* reflect "the knowledge of the future-predictors" or are "the untapped opportunities available to future-creators"? On one hand, "Opportunity Cost" claims that prices can drive incentives and therefore probabilities, but on the other hand, "Market Efficiency" claims that probabilities must in turn drive the prices. What will happen when these seemingly-contradictory forces collide?

Market Efficiency will always prevail. While *some* individuals may have *some* preference about the Vitcim (that he die, that he live, etc), or some ability to influence the Victim (to cause his death, or prevent his death, etc) *all* individuals prefer "having more money". These markets (both by the Assassin Market definition and by the Truthcoin design) are only for publically available events (remember that the entire marginal advantage to assassinating someone with a PM [as opposed to just privately contracting an assassination] was to get the public's money), ensuring that a money-loving public knows about the issues.

Don't Take My Word For It

Academic research has shown, both theoretically[8] and empirically[9], that probabilities out-drive all other forces.

Can we point to real-world instances where AMs were seriously tried, and yet they didn't work? I think so.

While it has never before been possible to create public financial markets *which are labeled* "Assassination Markets", markets which pay out a reward based on a certain individual's death (so, in consequence, *are* assassination markets) have always existed. There is already plenty of money for an Assassin to make in this way. It doesn't exactly take a Finance PhD to understand that the assassination of some President, Congressmen, Banker, CEO, etc would have a predictable and

[8] http://www.fhi.ox.ac.uk/manipulator-can-aid-prediction-market-accuracy.pdf

[9] http://mason.gmu.edu/~rhanson/biastest.pdf

exploitable effect[10] on certain assets (equities, gold, volatility-linked derivatives etc). With the trivial borrowing and leverage abilities of the average person, the trades could net quite a bit of money.

Someone could short Ahold (AMS:AH) and then go into Stop & Shop and try to secretly spray poison on all the produce, in an effort to reduce the AH stock price. Bill Gates could short Microsoft and remotely command his private jet to crash into the headquarters building…it isn't that these phenomena are fundamentally impossible, just that there are many more things *preventing* them from happening, than there are *encouraging* them to happen. People who are rich / sophisticated enough to have the working capital required to cash out the Assassin bounty usually have too much to lose by going to prison, and too many other (more-reliable) ways of achieving their political goals.

PMs Would Need Help to Become AMs

Without my special SI modification, PMs simply will not operate in this undesirable way. PMs ***react to*** the real world, not the other way around. After all, *some* people may want *some* particular thing to happen, but *all* people prefer having more money to having less money.

So far, my defense of PMs may seem trivial. After all, can we not simply add some assassination-themed SIs?

[10] http://www.usatoday.com/story/money/markets/2013/11/21/stock-market-reaction-to-jfk-assassination/3662171/

3. How the T-DAC Works

Intro

First of all, there's nothing suggesting it *will* work at all. The T-DAC and SI's are completely new and theoretical concepts invented by me. They have not been tested by anyone, they have not even been peer-reviewed. They assume that the SI's can dampen the circularity problem to such a degree that only "the one who can provide the good" (and set the SI) has positive expected value on his trades, and that any front-runners (who can't set the SI) have a sufficiently-negative expected value that they will not bother.

I will first explain how the SI's work, to show that it is a requirement that they be **publicly** editable *only* by **the provider** of the public good. This requirement makes them unfit for any anonymous activity ("Crimes").

The Schelling Indicator Disarms the Circularity Problem

The SI, in theory, separates *forecasts* about the act from *ownership proofs* of the act. With SIs, it is possible to successfully *forecast* that someone will build something, yet fail to "construct a proof that you were the individual who built it", and it is possible to "successfully prove that you built something" without *affecting a forecast* about the building of that thing.

Crucially, the SI can be updated by the project's **owner**. The info-timeline is "project work" -> "project has an owner" -> "owner makes a project-related selection of SI". The **order of actions** taken by the **owner** (the "Assassin" in the AM case) is of great significance and is as follows: [1] "observe market", [2] "create trade", [3] "create good (with SI)".

The SI Exists to Prevent Front-Running

By "front running", I refer explicitly to two types: [1] "slow" front-running (the circularity problem, where someone trades because they expect a large trade in the future), as well as [2] “fast” front-running (attempts to look at a single trade-in-progress and “run in front” of it). “Fast” front-running includes bribing miners (directly, or indirectly with fees) to include-or-not-include a transaction, or try to orphan/include a specific block.

How does this work? Simply: if the trade *is* front-run (in any way), the owner can just *change* the SI to something else, and make a new trade. Instead of draping the Lighthouse with a red flag, he will drape it with a blue flag. Changing the SI costs the owner nothing (as, at this point, he hasn’t even started working on the project); in contrast, the shares purchased by the front-runner are guaranteed to be worthless (dissuading anyone from attempting such a maneuver in the first place).

The SI Always Leaks Identity (or Timing, or Both)

The SI can *only* prevent front-running **if the provider can modify it and no one else**. If any non-providers could modify the SI, they could themselves become front-runners and steal the contributions. Notice the huge difference here between public *goods* and *bads*. The (publicly known) owner of a public *good*, can easily modify the SI (by making a public statement, updating their website or product). He or she has plenty of time to do this, and can do it a theoretically unlimited number of times. The (anonymous) owner of the public *bad*, by contrast, has a big problem. The SI contains a proof-of-ownership, and yet with crimes the entire challenge is to *avoid* producing any evidence that you “own” the crime. The owner of a public bad would likely use try something like “assassination to occur on an {odd/even} numbered month” or the days of the week (notice, we are approaching the original AM proposal).

This reintroduction of private “Assassin-only knowledge” may feel like it is getting close to something which will work. Unfortunately, the exact function of PMs is to take private information and make it public, which they will do in this case as soon as the Assassin claims his money (which is *before* the murder takes place). In short: while the Asssassin can “set” the SI, he cannot “reset” it.

One can't murder someone on Wednesday, un-murder them on Thursday and then re-murder them on Friday (if the trade for "Wednesday" was front-run in any way, or it wasn't accepted into a block, or it's block was orphaned, etc). In fact, this 'finality' generalizes well to "all crimes". If "crimes" could easily be un-done, we wouldn't really need to worry about them. We'd just undo them.

Perhaps the SI will involve some kind of cryptographic layer, such that the criminal can [1] perform the crime, [2] link proof of the event with an individual public key, and [3] cycle through the different proofs (by signing updates), such that only the criminal has the last word.

In this case, the concept of linking the SI to the good's owner has not been solved, it has merely been obfuscated with cryptographic labels. How, exactly, would one "link proof of the event to a single public key"? The issue is not the post-event linkage. While it is clear that, the Assassin may, at some inconvenience (and literal evidence of guilt), take HD video of the Victim's body alongside a QR code. The issue is the pre-event linkage: the SIs are all known in advance! Public keys are only useful if only the Assassin knows the corresponding private key. Moreover, the Victim can himself pre-record faked versions of similar death-videos (and the resulting he-said-she-said game favors the Victim). Fundamentally, **this information is either "only known to the Assassin" or it isn't.** And the Assassin needs this information to be revealed (via their trade) to predate the final proof of the crime by the settlement time (an hour).

This will not be a problem with *non-crimes*. Recall that, in the public good sequence, the trade comes first, so if a front-runner buys all of the SI-states, the provider can simply choose not to trade, not to set the SI, and/or not to provide the good at all (at which point the shares purchased – albeit cheaply– by the front-runner are completely worthless).

There is no way to escape this fundamental time-constraint. For example, consider a seemingly viable alternative where one constructs the proofs before the crime is publicly observable, for example by secretly slipping the Victim some very slow acting poison, and constructing proofs of this deadly dead while the person is still healthy, trading on these proofs, and finally cycling the proofs until they land. Setting aside the fact that the Assassin has just *revealed* to the Victim that he has been slow-acting poisoned (and he can now rush to the hospital), how exactly can the Assassin prove that the Victim has been given a *deadly* dose of poison, before the Victim actually dies? How can you prove that the Victim has actually been given *any* quantity of poison at all, or that *what was given* was actually poison? It is either a proof ***of the event*** or it isn't.

While the individual SI-states can refer to private information, the trade purchasing the SI-shares is known to the whole blockchain network, and therefore always public and can therefore always be front-run.[11] Even if the SI could be somehow reset a few times, each of the SI-trades could be front-run. The public good provider can reset as many times as he likes, but to the public bad provider this is very costly (if not impossible).

4. Crime in the Real World

A Criminal Contradiction: Desperation vs. Working Capital

Yet another distinction between AMs and PMs: AMs pay if you guess the death-date correctly, PMs pay if you can *secretly* place *huge amounts* of money *at risk* on a SI that you think (hope) you can influence (without anyone learning of your plan or of the trade). Criminals are not known for their access to working capital. Only a few wealthy individuals can do this without some kind of loan, and someone with this access to cash is usually not desperate enough to absolutely need to have a heinous crime performed immediately.

[11] Truthcoin has anti-front-running measures, but the costs imposed by the measures are negligible for a single pre-prepared trade. Similarly, someone concerned for their life might pay mining pool operators to not include such a transaction in their blocks, or try to orphan one should it appear in a block. Furthermore, mining pool operators could simply censor these transactions for personal ethical reasons. Fungibility is an essential property of money, but not of markets.

Perhaps the Assassin might take out a loan. However, loans are a written agreement…a legal contract, at which point the Assassin runs into the exact same contract-concerns I expressed earlier. Loans generate records, and a paper trail. There will always be a rich person who needs to know what the money is for; this individual would be held accountable both by the law (going to prison if the plot is uncovered or ratted out) and the free market (by losing their money if something goes wrong with the contract).

Perhaps a rich person (not needing a loan), newly motivated by this smart contract, hires a poorer person to carry out the deed. We would now have a desperate Assassin and a more-insulated Assassin-employer. This changes nothing. After all, there is still a rich person who can go to jail (and who can be ratted out). Depending on the payment terms, the poorer assassin might keep the money without doing the assassination, or blackmail the richer person from that point forward. The only real difference expressed here is the transition from individual to team. This has the same "contract" problem, mentioned earlier, that teams are difficult to sustain in illegal activity because contracts can't be enforced (and are even negative-enforced).

Readers may wonder if the same difficulties are present in the public goods (non-crime) case. Quite the reverse: this problem is avoided entirely. Those providing public goods can prove how much they will receive if they successfully produce the good. After the good has been produced, the loan to finance the single large trade claiming the prize is virtually risk free. Your borrowing costs are low. In fact, because they are known to be low, the good-provider can safely delay taking the loan out until it is immediately necessary (lowering the duration of the loan, which lowers the total loan cost further).

The Worst Case ("The Easiest Crime")

Let us broaden the scope of our vigil, as wide as possible: from "assassination" to "anything dishonest". It would seem, that one final challenge has been overlooked: cases where both [1] the provider of the crime-service is just one person (ie, no need for potentially-incriminating, non-enforceable written contracts), and the provider [2] can provide the service for free (no marginal effort is needed on his part, and no marginal suspicion accrues to him specifically). Strong examples would be the cases of legal arbitration and sentencing (for example, "Will Supreme Court Justice Anthony Kennedy (the key swing vote) rule in favor of gay marriage during June 2015?"). The concern here is not that any individual will Contribute to alter the outcome (which I've previously described as impractical and self-defeating), but that the mere existence of the PM may alter an individual's decision making process.

The Problem

Given that sentencing can be influenced by random rolls of dice, the obvious concern would be that a judge (or juror) would turn to an existing PM to lighten their cognitive load. If in an agreeable mood, they may outsource their decision to the PM; if in a contrarian or spontaneous mood, they may oppose the expectations of the PM. Of course, this is true of all forecasts (and all speculation), but in a hypothetical future world, where all share the author's opinion that PMs are categorically the optimal forecast (in addition to being easy to read and understand), the problem may worsen.

Moreover, the PM will be offering market prices which are based on the publicly known *objective* facts of the case (and the known biases of the judge/jurors). This, paradoxically, presents "judges and jurors with the option to trade in these markets" with more money-making opportunities if they *ignore the objective facts as much as possible*. The PM is, in a sense, eternally trying to *bribe* the judge or juror to submit (what it feels is) the non-objective ruling. This problem is small: by definition, these single-party rulings are inherently subjective, and it is kind of bizarre to imagine the forecast drifting too far away from the (incentive-neutral) 50%-50% prices (and, in practice, the usual challenge with PMs is that public interest in them is too low). Nonetheless, this "diminished objectivity incentive" is something to be aware of.

Mitigating Factors

For this "crime" of bribery/corruption (if a PM is made on judicial topics), we (both law enforcement and the public) have access to supervisory resources which we typically lack: we know that a "bribe" has been offered, we know the amount of the bribe, we know who to whom it has been offered, and we know the desired corrupt action. We even know something about the status of the bribe (we might think it has been "accepted" if a large trade suddenly swings the odds, and "rejected" otherwise). Police can't do some things (prevent the bribe-offer), but they can do plenty of other things (watch the bribe and the individuals involved).

Secondly, *we already live in a world* where certain court cases (for example gun control cases), have a profound effect on certain markets (stock price of weapons manufacturers). Although the world of the Blockchain does offer more financial privacy and convenience, it also offers permanent and transparent blockchain-evidence, and cryptographic evidence of guilt.

Solutions

Ultimately, it seems, items with "high bribes" (high profile court cases, on which a lot of money is wagered) would warrant correspondingly high scrutiny. Official policies could (and should) dictate that extra steps be taken if PMs are created on bribable, individual, subjective decisions of great importance (an obvious extra step would be to introduce peer review [rival judges examining the process for any signs of corruption]), and penalties for "accepting" such bribes should be correspondingly severe (US law should vehemently punish judicial decision-makers who trade in any PMs on their own rulings).

5. Superior Alternatives

Let's examine the very motivation behind the creation of the AM in the first place, to see if the original need can be better met: Jim Bell's original essay objected to an unrepresentative leadership, and hostile to the very concept of hierarchy (in which the worst would rise to the top, where they could do the most damage). How else might the masses use crypto-economics (under the AM assumptions) to check the power of an entrenched leader?

A Tame (But Still Superior) Alternative

The first obvious answer is that one might **pay the person to *resign*** (and be replaced by someone else). I mean does the AM-user really *need* the individual in question to *die*? That would seem to be unnecessarily extreme.

The resign-market is superior to its AM counterpart. If two hypothetical T-DACs would produce the same effect (getting some guy out of your hair), then it is reasonable to believe that Market Authors would prefer to make *only* the one which most avoids problems with [1] the law, [2] Branch-rules, or [3] trader ethics. Even an inherently violence-preferring person might more-prefer to join forces with a larger (nonviolent-preferring, or at least violence-indifferent) group.

Such a payment the negative version of campaigning to get a certain individual hired. It is roughly as ethically ambiguous as the use of NDAs / non-competes to prevent employee poaching (not very unethical at all).

The Cheaper, Safer, More-Effective Option

So far, we've had people buying assassinations in a (presumably expensive) store. But why *spend* money (to get what you want) when you can *make* money (getting the exact same thing)? Activism is more fun when *they* pay *you*.

Markets let traders express their opinions in a constructive way. Instead of plotting murder, a Market author can plan winning Markets and winning trades. Using conditional bets, individuals can wager that, if Politician X is elected (or re-elected), something that they care about will harmed. If these users are right, they will *make* money; meanwhile, everyone else will be forced to listen to what their money has to say!

Users can trade in Market-topics on every decision that they make –every law that is passed or signed–, on a case by case basis. In a way, these individuals are running the organization (or country) now. Either Politician X does a good job, or *the user profits individually from X's refusal to act on the user's profound wisdom*.[12]

6. The Anti-AM Trump Card

The "good" financed in an Assassination Market is the Victim's death. Yet, the Victim is in arguably the best position to cause this death. In fact, the Victim will almost always be able to *fake* his own death or disappearance.[13] In this way the Victim can collect the death-bounty from his enemies (the Contributors), then reappear after the PM resolves. (If the PM is made to not-resolve for a long time, this is inconvenient to any 'real' Assassins…not *as* inconvenient , but a disincentive to perform the task, and nonetheless an incentive for Authors to keep the resolution soon.)

Obviously this is extraordinarily inconvenient to the Victim (who is still largely "a victim"). Such an elaborate and serious deception might be disastrous to the Victim's relationships (more so the longer the person must remain "dead"). A spouse could remarry, children's graduations could be missed, crucial disruption to the Victim's career (as well as to any organizations of which he is a member), life insurance / inheritance would need to be rewound, etc.

[12] Or, of course, your beliefs were wrong the whole time, in which case you more-than-deserve your financial losses.

[13] Law enforcement may even help you do this.

However, this *option* is worth discussing, as it highlights fortunate shortcomings of AMs, relative to traditional murder-for-hire. One shortcoming to AMs is that the Victim is forewarned: these "market activities" are only effective if they are public. A second is that the Contributors must pay up front, with no *guarantee* that their payments won't go to a Victim who has merely faked his own death. Perhaps this "faking your death" option, even if expensive to use, might significantly dissuade murders from trying to use AMs at all (let alone PMs-as-AMs). Individuals who only wish to rid themselves of an individual temporarily, or in a certain context (at work, as President), can use PMs to encourage those steps in a way that is entirely more legal/ethical (as described above), and also more likely to work.

7. Amplifying the Disadvantages

I have demonstrated a core set of difficulties present in using PMs to encourage crimes. With this "skeleton" in place, Truthcoin's design allows it to leverage this skeleton to further resist crime-encouragement.

Truthcoin is censorship-resistant, not censorship-proof. It explicitly censors questions ("Decisions") that are too confusing or off-topic. Branches (to which each Decision must belong) can, at first, set their own guidelines for determining if something is "off topic", although these initial guidelines are more difficult to change later. New Branches cannot be created without permission from an existing Branch, and Branches which go unused "die" and exit the system.

With this in mind, I plan to create the first Branch with a guideline expressing that anything violent[14] is defined as 'off topic', thus completely censoring violence from the initial Truthcoin system.[15]

[14] I also intend to ban anything involving a country's judicial system, but I'm not sure if this will last. The judicial system already affects existing financial markets, so I'm not sure to what extent this would be a problem.

Of course, the rules might be changed, either [1] explicitly on the original Branch or [2] on future child-Branches. However, in light of the existing disadvantages, one would hope that a critical mass of VoteCoin-owners would find such a rule-change to be pointless at best and risky at worst. Likely only a small minority (and/or, those individuals who intend to harm the protocol's reputation), would be interested in changing the rules of their Branch in this way. With the persuasion route unlikely to succeed, malicious agents much "purchase the corporation" (the Branch) to push their change through. Of course, even this (expensive) strategy may prove to be ineffective: if everyone moves to a competing Branch, the change may result in hazardously low Decisions/fees, such that the economic security of the Branch falls to a level where honest reporting is no longer likely. This would lead the Branch will spiral downward into disuse and eventual "death", and the attacker's "investment" will go to zero (having achieved nothing).

8. Nirvana Fallacy

Today's world is more violent than most would prefer. The past was violent, the near future will probably be violent. Perhaps the period of greatest violence will be in the future, thousands of years from now, when centuries of peace are disrupted suddenly by the instantaneous obliteration of trillions of lives (extraterrestrials, space travel disaster, deranged physicist, planetary WMDs, etc). Who can say? The relevant question is: how can we organize society today to reduce expected total violence?

With that in mind, the current analysis is woefully incomplete, as it fails to take into account any assassinations *prevented* by PMs. As stated above (section XX!), PMs have the potential to pave the way for greater civilian-politician trust and a more civilized political process: some anxious citizen might be, today, planning desperately to assassinate a politician, because they fear

[15] Although PMs on violent subjects would almost certainly be able to prevent real-world violence (terrorism deaths, troop casualties) it is obviously best not to impose such an idealistic modification to the safety of *other* people until one is absolutely certain that all of the relevant facts and viewpoints have been considered.

oppression or feel that the peoples' values are being ignored. Put simply, PMs allow information to be expressed, and then force it to be respected (as prices are what game theorists call "common knowledge"). In this way, they help "fringe knowledge" (conspiracy theories, emotionally-disturbing truths, info-coercive activities, etc) get noticed (if it is credible), and –because they do this- PMs also can put individuals at ease that no relevant fringe knowledge is lurking unnoticed. With more transparency imposed on the governance process, there is less of a reason for individuals to feel individually responsible for a drastic maneuver to "free" the "brainwashed" masses.

As mentioned earlier in this document, PMs generally have the potential to increase economic growth, on firm and national scales, through vastly improved governance. By *failing* to seize an opportunity to improve governance, we impose needless harm on society. It is unknown how many lives are lost every year as a result, but it may be on the order of millions.