

Risk Assurance for Hedge Funds Using Zero Knowledge Proofs

Michael Szydlo

mike@szydlo.com

Abstract. This work introduces a new tool for a fund manager to verifiably communicate portfolio risk characteristics to an investor. We address the classic dilemma: *How can an investor and fund manager build trust when the two party's interests are not aligned?* In addition to high returns, a savvy investor would like a fund's composition to reflect his own risk preferences. Hedge funds, on the other hand, seek high returns (and commissions) by exploiting arbitrage opportunities and keeping them secret. The nature and amount of risk present in these highly secretive portfolios and hedging strategies are certainly not transparent to the investor.

This work describes how to apply standard tools of cryptographic *commitments* and *zero-knowledge proofs*, to financial engineering. The idea is to have the fund manager describe the portfolio contents indirectly by specifying the asset quantities with cryptographic commitments. Without de-committing the portfolio composition, the manager can use zero knowledge proofs to reveal chosen features to investors - such as the portfolio's approximate sector allocation, risk factor sensitivities, or its future value under a hypothetical scenario.

The investor can verify that the revealed portfolio features are consistent with the committed portfolio, thus obtaining strong assurance of their correctness - any dishonest portfolio commitment would later serve as clear-cut evidence of fraud. The result is a closer alignment of the manager's and investor's interests: the investor can monitor the fund's risk characteristics, and the fund manager can proceed without leaking the exact security composition to competitors.

Keywords: Hedge fund, zero-knowledge, commitment scheme, investor trust.

1 Introduction

This paper describes a novel application of zero-knowledge techniques to the relationship between an investor and a portfolio manager. The interest of the fund manager is in earning high returns, so he may want to keep his exact portfolio and trading strategy secret. An investor, on the other hand, also requires mechanisms to ensure the honesty of the managers, and to check that the fund's risk characteristics are in line with his own risk preferences. We address the fundamental problem of how to control the flow of risk information to serve these

distinct interests. We suggest that the tool described in this paper is particularly suited to *hedge funds*, which tend to be highly secretive, more loosely regulated, and potentially very lucrative.

Cryptography has been applied to financial transactions before, in both generic ways (privacy, integrity), as well as in ways specific to transactions (digital cash, and privacy-preserving auctions). Rather than focus on the transactions themselves, our approach uses cryptography to allow a more finely controlled release of financial information to an investor.

Our idea is to use cryptographic commitments and zero knowledge proofs in a remarkably simple way: The fund manager describes the portfolio contents indirectly by specifying the asset quantities with cryptographic commitments. Then, without de-committing the portfolio composition, the manager can use zero knowledge proofs to reveal chosen features to the investor. This technique differs from traditional topics in financial cryptography, since it applies the tools of cryptography directly to mainstream *financial engineering*.

The main cryptographic tools we require are standard: *Pedersen Commitments* and *Interval Proofs*. We review the mechanics of these tools and show how to assemble them into (zero knowledge) statements which are meaningful to the investor. We stick to such well-known building blocks in this paper in order to retain the focus on the new finance application.

We have implemented a prototype of the protocol to demonstrate its feasibility. Despite the potential for efficiency improvements, the basic construction is already good enough to serve in practice. This shows that it is possible for a fund to communicate interesting risk information for large and complicated portfolios on a daily basis.

The rest of this paper is organized as follows: In Section 2 we provide some background on hedge funds, and the risks associated to them. In Section 3 we review the cryptographic building blocks we require, and in Section 4 we describe the mechanics of the protocol. We continue with some detailed applications in Section 5. In Section 6 we describe the results of our prototype implementation, and discuss efficiency concerns. We conclude in Section 7, and provide an appendix with some further technical details on the cryptographic construction.

2 Finance Background

For the non-finance professional, and to motivate our work, we first review some basic finance material, highlighting the roles of information and risk. We focus on the differing interests of the investor and fund manager with respect to release of information to motivate the need for our risk communication protocol. Including the background on common methods employed in the industry to measure risk also helps show that most of the meaningful risk statements used in practice are compatible with our protocol. Much of this material is present in introductory finance textbooks, e.g., see [5], which emphasize quantitative methods.

2.1 Hedge Funds and Risk

Portfolios and Risk: An investment portfolio is just a collection of assets designed to store or increase wealth. In a *managed fund*, the *investor* turns over capital to a *fund manager*, an investment professional who buys, sells, and otherwise maintains the portfolio in return for a fee or commission. The assets often contain publicly traded *securities* such as stocks, bonds, commodities, options, currency exchange agreements, mortgages, “derivative” instruments, as well as less liquid assets such as real estate, or collectibles. Examples of managed funds are pension funds, 401K plans, mutual funds, and hedge funds.

Every type of investment contains uncertainty and risk. Ultimately, the risk inherent in investments derives from the fact that the future market value¹ depends on information which is not available: information concerning either unknown future events, or information concerning past events which has not been publicly disclosed or effectively analyzed. The charter of the fund manager is to manage these risks in accordance with the preferences of the investor.

Risk Factors: The finance profession has developed a plethora of models to define and estimate portfolio risks. A first description of a portfolio’s risks includes a breakdown of the types of assets in the fund such as the proportion of capital invested in equity, debt, foreign currency, derivatives, and real estate. A further breakdown specifies the allocation by industry type or *sector*, or region for foreign investments.

The future value of an investment depends on such future unknown factors as corporate earnings for stocks, interest rates and default likelihood for bonds, monetary policy and the balance of trade for foreign currency, regional political stability for any foreign investment, re-financing rates for securitized mortgages, housing demand for real estate, etc.

Risk models identify such measurable *risk factors*, and study the dependence of the asset’s value on each such factor. Such *factor exposures*, are estimated with statistical regression techniques, and describe not only the sensitivity to the factor but also how the variance, or *volatility* of a security depends on such correlated factors. Assembling such analysis for all securities in a portfolio, the fund manager has a method for quantitatively understanding the relative importance of the risk factors his portfolio is exposed to. Another important tool, *scenario analysis* estimates the future value of a portfolio under a broad range of hypothetical situations.

Hedge Funds: To *hedge* against a risk is to effectively buy some insurance against an adversarial event. When two assets depend oppositely on the same risk factor, the combined value of the pair is less sensitive to that factor. A *Hedge Fund* is just a type of portfolio designed to have certain aggregate risk characteristics. Hedge funds may use leveraging techniques such as *statistical*

¹ Economists like to point out that there is no robust intrinsic definition of value outside a market.

arbitrage, engaging in long and short positions in similarly behaving securities, hoping to earn a profit regardless of how the correlated securities behave.

Hedge funds are often large private investments and are more loosely regulated than publicly offered funds. (Only in 2006 must hedge funds register with the SEC at all). Such extra flexibility affords the possibility of exceeding the performance of more standard funds. For example, hedge funds often take a position contrary to the market consensus, effectively betting that a certain event will happen. When accompanied by superior information or analysis such bets can indeed have high expected value. Of course, highly leveraged funds can be extremely sensitive to a particular risk factor, and are thus also susceptible to extreme losses.

The high investment minimums, lax regulation and secrecy or “black box” nature of hedge funds has fostered an aura of fame and notoriety through their spectacular returns, spectacular losses, and opportunities for abuse. Recently, though, there has been interest in marketing hedge funds as viable opportunities for the average investor.

2.2 The Role of Information

Information and Asset Prices: A *market* assigns a value to an asset based on the prices in a steady stream of transactions. It is the pieces of information which are perceived to be relevant to the asset’s value which are compared to existing expectations and drive the supply, demand, and market price. The pivotal role of information is embodied in the *efficient market hypothesis* which states that under the assumption of perfect information distribution, the collective brainpower of investors will reduce arbitrage opportunities, and force the market price to an equilibrium.

In the real world, information distribution is not perfect, and the *information asymmetries* among parties significantly affect the behavior of asset prices in the market. The situation is worse for illiquid assets, for which one must rely on some ad-hoc *fundamental analysis* to estimate the value. Similarly, it is difficult to assign a robust value to an investment fund with opaque risk characteristics (such as a hedge fund). An increasing sharing of the actual risk profile of hedge funds would increase their usefulness in *funds of funds*, for example.

The Importance of Secrets: Certain investments, such as passive funds which track an index may have no requirement to protect the portfolio contents or trading patterns. Actively traded funds, on the other hand, have good reasons to maintain secrets. For example, revealing in advance an intention to purchase a large quantity of some security would drive the price up. A parallel can be made with corporations: Sharing technological, financial, and trade secrets would undermine the competitive advantage of a firm.

Especially relevant to our focus, if a hedge fund were exploiting a subtle but profitable arbitrage opportunity, revealing this strategy would quickly

destroy the benefit, as other funds would copy the strategy until it was no longer profitable. Thus, a rational investor will support such constructive use of secrets.

The Importance of Transparency: Secrecy is also dangerous. The actions of a fund manager might not always represent the goal of creating value for the investor! The danger of too much secrecy is that it also reduces barriers to theft, fraud, and other conflicts of interest. An example of corrupt behavior that might be discouraged by increased transparency is the practice of engaging in unnecessary trading motivated by brokerage commissions. To combat this risk, individual investors require enough access to information about a company or fund to help ensure honest management, consistent with the creation of value.

Another kind of problem will arise if the investor is not aware of the kinds of risks his portfolio is exposed to. In this case it is impossible to tell if these risks are in line with his preferences. A fund manager might be motivated by a fee structure which encourages him to take risks that are not acceptable to the investor. When the fee structure or actual level of risk in the portfolio is not evident to the investor, a fund manager may legally pursue actions consistent with interests other than the investor's.

Aligning Interests: The above discussion about the differing views concerning just how much risk information should be kept secret and how much should be revealed shows how difficult it is in practice to perfectly align the interests of investors and fund managers. The traditional approaches to mitigating this problem involve financial regulatory bodies such as the SEC, which seeks to institute reporting laws and support capital requirements that protect the investor, ideally without imposing too large a burden on the financial institution. In the case of hedge funds, the position of the SEC is that the interests of the investor are not adequately protected [1]. Indeed, it has not been able to eliminate all fraud and conflict of interests arising in the context of hedge funds.

There are several requirements for a good set of mechanisms to align the interests of investors and managers. These include methods for the investor to ensure the honesty of the fund manager, methods for the investor to be aware of the fund's evolving risks, and contractual agreements and fee structures which discourage the manager from adding hidden risks. Finally, the mechanisms should not discourage the fund manager from fully exploiting any competitive advantage or superior analysis which he might have.

2.3 Finance and Cryptography

Previous Work: There are many existing applications of cryptography to financial infrastructure. The most significant practical applications involve well known aspects of securing the transactions themselves: providing authenticity of the parties, integrity and non-repudiation of the transactions, and confidentiality among the parties. Such applications all use cryptography in a generic way, not tailored to any particular requirements of finance.

More interesting advanced finance-related applications of cryptography include fair exchange, secure auctions, and digital anonymous cash. These appli-

cations use cryptography as a building block to compose cryptographic protocols which protect some aspect of a transaction, preserving some secret, or prove the correctness of a protocol step. The technique of sending non-interactive proofs relative to previously committed values is pervasive in protocol design.

The present application to finance is not directly focused on the transactions, but instead on the release of information about the evolving portfolio's composition and risks. This kind of application has not previously appeared.

New Contributions: Our contribution is the proposal of an additional mechanism which will help achieve a better balance of information sharing between fund managers and investors. We present a protocol which can precisely control the level of transparency in an investment fund. The result is that the investor can ensure that an appropriate level and type of risk is taken, yet the fund can pursue competitive strategies which would not be possible if the restriction of perfect transparency were imposed.

Cryptographic commitments, and zero knowledge proofs provide versatile tools for precisely controlling the delivery of partial and verifiable pieces of information. Our work is the first to exploit these methods in the context of financial risk management. When our protocol is used to communicate the amounts and types of risk in a portfolio, the interests of each party will be better served. In addition to outlining the basic approach, the technical applications we describe below serve as specific examples of how various types of risks can be communicated within our framework.

3 Cryptographic Building Blocks

The cryptographic tools we require in our construction are all standard. Namely we require commitments with a homomorphic property, and zero knowledge proofs that a committed integer lies in a interval. In this section, we review the most well-known versions of these constructions. Throughout this paper, we let p denote a large prime and q a prime such that $q|p-1$. Let $\mathbf{G} = \mathbf{Z}_p$ denote the group of mod- p integers, and let $g \in \mathbf{G}$ and $h \in \mathbf{G}$ be group elements of order q such that the discrete log, $\log_g(h)$ is unknown. We also let hash denote a cryptographic hash function with range $[0, q-1]$.

Pedersen Commitment: A cryptographic commitment is a piece of data which binds its creator to a unique value, yet appears random until it is de-committed. A *Pedersen commitment* [8] to x with randomness r is the group element $C_r(x) = g^x h^r$, and can be de-committed by revealing the r and x . This commitment is computationally binding and unconditionally hiding. Since a commitment can only feasibly de-commit to the original value of x , we also say $C_r(x)$ "corresponds" to x .

Linearity Property: We make essential use of the linear (homomorphic) properties which Pedersen commitments enjoy:

$$C_r(x)^a = C_{ar}(ax) \tag{1}$$

$$C_r(x)C_{r'}(x') = C_{r+r'}(x + x') \quad (2)$$

Thus, without knowing the values x and x' that two commitments hide, any party can compute a commitment to any fixed linear combination of x and x' .

Proof of Knowledge: A *zero knowledge proof of knowledge* allows a prover to demonstrate knowledge of hidden values without actually revealing them. A proof of knowledge of a (Pedersen) committed integer x [10] demonstrates knowledge of some x and r such that $C_r(x) = g^x h^r$. We focus on *non-interactive* proofs of knowledge, for which the proof is concentrated in a single piece of data and can be later verified without any further participation of the prover.

One can also prove that a committed value x satisfies some condition $\phi(x)$ without revealing it, and we use the notation $POK(x, r \mid C = g^x h^r, \phi(x))$ to denote a zero knowledge proof of knowledge of (x, r) satisfying both $C = g^x h^r$ and the predicate $\phi(x)$.

Schnorr OR Proofs: The well known *Schnorr OR proof* [6, 10].

$$POK(x, r \mid C = g^x h^r, x \in \{0, 1\}) \quad (3)$$

can be used to prove that $x \in \{0, 1\}$, (provided this is true), without leaking whether x is 0 or 1. The proof data consists of the five values $\{C, r_1, r_2, c_1, c_2\}$ such that $c_1 + c_2 = \text{hash}(a_1, a_2) \pmod{q}$, where $a_1 = h^{r_1} C^{-c_1}$, and $a_2 = h^{r_2} (C/g)^{-c_2}$. Any verifier can efficiently check these conditions. In Appendix A, we review the *completeness*, *zero-knowledge*, and *soundness* properties of this construction.

Interval Proofs: We will need proofs that a committed integer satisfies an inequality such as $x \geq A$. One way to accomplish this is to prove that x lies in an interval $[A, B]$ for a large enough B . We now review the classic interval proof [4, 7, 6], based on bounding the bit length of an integer.

$$POK(x, r \mid C = g^x h^r, x \in [0, 2^k - 1]). \quad (4)$$

The proof is constructed as follows: First expand x in binary: $x = \sum_0^k 2^i a_i$, and produce a commitment $C_i = C_{r_i}(a_i)$ for each digit. The commitment to the last digit is set to be $C/\Pi_1^k(C_i^{2^i})$, so that the relation $C = \Pi_0^k(C_i^{2^i})$ holds². Finally, for each digit a_i compute a Schnorr OR proof demonstrating that $a_i \in \{0, 1\}$. This proof is verified by checking the list of k Schnorr proofs, and checking that $C = \Pi_0^k(C_i^{2^i})$ holds.

To construct a proof that x is in the range $[A, 2^k - 1 + A]$, one simply follows the same procedure, replacing C with C/g^A . These proofs are reasonably efficient in practice, as long as the interval is not too large. See [3] for alternate constructions of interval proofs designed for time and space efficiency.

² An alternative to adjusting the last digit's commitment is to add a proof that C and $\sum_0^k 2^k C_i$ commit to the same number.

3.1 Further Notation

For our application we will need to make commitments to a large set of quantities (assets) and prove statements about linear combinations of them. We consider a universe of asset types $\{A_i\}$, and let b_i denote an amount of asset type A_i , and C_i a commitment to this value.

By virtue of the homomorphic property of Pedersen commitments, for any list of coefficients $\{m_i\}$, the product $\prod C_i^{m_i}$ is a commitment to $\sum m_i b_i$, and can thus be publicly computed from the $\{C_i\}$ and $\{m_i\}$. By using the interval proof technique reviewed above, the creator of the commitments can prove that $\sum m_i b_i \in [Q, Q + 2^k - 1]$, for any threshold integer Q . Since all of the zero-knowledge proofs we use are with respect to the same C_i , hiding b_i we abbreviate

$$POK(x, r \mid \sum m_i C_i = g^x h^r, x \in [Q, Q + 2^k - 1]) \tag{5}$$

to the more succinct expression which also de-emphasizes the interval length

$$ZKP_k(\sum m_i b_i \geq Q). \tag{6}$$

Similarly, a zero knowledge proof that an expression is bounded above is denoted $ZKP_k(\sum m_i b_i \leq Q)$. To summarize, this proof data (6) allows any verifier with the $\{C_i\}$, $\{m_i\}$ and Q to check that $\sum m_i b_i \geq Q$ for the b_i hidden in the C_i .

4 The Risk-Characteristic Protocol

4.1 The Basic Approach

The process we describe provides the investor with a new tool to verify claims made by the fund manager, and there are both contractual and cryptographic aspects of the mechanism. Additionally, the involvement of a third party enhances the effectiveness of the scheme.

As part of the financial design phase, a universe of possible asset types is chosen, and the kinds of risk information to be verifiably communicated are identified. Such parameters are incorporated into the contract governing the fund. The more interactive component of the scheme involves a periodic delivery of risk assertions and accompanying proofs to the investor.

Contractual Aspects: The legal document governing the investment, the *prospectus* specifies the rights and obligations of the investor and the fund, including the mechanics of the contributions, payments, withdrawals, and fees. The prospectus may also specify or limit the types of investments made within the fund.

With our scheme, the architect of the fund chooses the risk profile and management strategy that he will follow, and incorporates the investment restrictions he is willing to guarantee into the prospectus. As part of a legal agreement, the fund would already be legally obligated to respect these conditions. However, such guarantees become much more meaningful when there is a mechanism for the investor to verify them in real time. The following steps facilitate this.

Within the prospectus a list of *allowable assets* is specified. The assets A_i can be directly identified by symbol if the security is market traded, and if not, described via their characteristics. Illiquid or private assets such as real estate, commercial mortgages, private bonds, or reinsurance contracts, can still be identified by descriptive categories. The units must be specified for each security, or asset type, since the rest of the protocol requires that the quantities be represented as integers. The *risk conditions* must also be expressed in the contract, and need to be expressed in a specific form to be compatible with the framework of our protocol. The conditions on the quantities b_i of assets A_i must take the form

$$\sum m_i b_i \leq Q \text{ or } \sum m_i b_i \geq Q \quad (7)$$

where the set of coefficients $\{m_i\}$ and bound Q determine the nature of the condition. We denote the list of conditions incorporated into the contract by Limit_j . It is easy to see how such conditions might be used to limit the amount invested in a single security, asset type, or sector.

In Section 5, we discuss how such conditions can also be used to bound total exposure to a specific risk factor, or expected value under a hypothetical scenario. Thus, the linear form of the conditions is not too restrictive. The applications using factor exposures or scenario analysis should also place additional data in the contract. The data which must be placed in the prospectus is thus:

1. The list of asset types A_i .
2. The list of conditions Limit_j .
3. (Optional) The list of risk factors F_j .
4. (Optional) The list of factor exposures $e_{i,j}$.
5. (Optional) The list of scenarios S_j .
6. (Optional) The list of scenario valuations $v_{i,j}$.

4.2 The Protocol Steps

Once the prospectus has been fully designed, the fund manager may solicit funds from investors and invest the capital in a manner consistent with the contractual restrictions. As often as specified in the contract, (e.g. daily), the fund manager will commit to the portfolio, and produce statements and proofs for each of the contractual risk-limitations. The commitments may also be sent to a third party to facilitate resolution of disputes. The protocol takes the following form:

1. The fund manager commits to b_i with C_i .
2. The fund manager delivers commitments $\{C_i\}$ to the investor, and optionally to a third party.
3. (Optional) The fund manager also sends a de-commitment of the committed quantities $\{b_i\}$ to the third party.
4. The fund manager asserts that conditions Limit_j are fulfilled, computes proofs $ZKP_k(\sum m_i b_i \leq Q)$, or $ZKP_k(\sum m_i b_i \geq Q)$, and sends them to the investor.
5. The investor verifies the completeness of the correctness of the proofs.

6. In case of dispute, the commitments may be opened or revealed by the third party. If the actual portfolio holdings do not match the committed holdings, the commitments serve as direct evidence of fraud.

We now elaborate on several aspects of this protocol.

Trading Behavior: In order to respect the contractual risk conditions, the fund manager must be sure to check that the risk profile would remain sound before effecting any transaction.

Commitment Step: Using the commitment scheme reviewed above, the number of units, b_i , of each A_i is committed to. The package of committed asset values is digitally signed and timestamped, and sent to the investor.

The commitments are binding - once made they can not be de-committed to a different value. This serves as a strong incentive against deliberate misstating of the portfolio. Of course, it is impossible to rule out the possibility that the fund manager lies about the asset quantities b_i in order to misrepresent the status of the fund. However, the quantity held of a particular asset at a given point in time is an objective piece of information. Making such a false statement would clearly be fraud.

Third Parties: We suggest the use of a third party to increase the effectiveness of the fund's incentive to commit honestly to the portfolio. For example, the committed portfolio might also be sent directly to the SEC, or to a different regulatory organization.

When the corresponding de-commitments are included in the message to the SEC, or other third party, this organization can also act as a trusted third party, confirming the correctness of the commitments, against independent information procured about the fund's contents, for example, by examining exchange records, and brokerage transactions. In this manifestation, the investor will have an even stronger guarantee, despite still never learning the asset quantities himself.

An alternative to the SEC would be another independent organization, such as a data storage firm, which would timestamp the commitment data, keep the de-commitments (if included) private, and readily provide the data to the court in case it is subpoenaed. If the protocol is implemented without sending the de-commitments to the third party, the commitments still serve as evidence should the court order them to be opened. A final option is to employ multiple third parties, and use the technique of secret splitting [11] so that two or more entities need to cooperate to obtain the data.

Computing the Proofs: The proofs of the form $ZKP_k(\sum m_i b_i \geq Q)$, $ZKP_k(\sum m_i b_i \leq Q)$ are computed according to the process reviewed in Section 3. One technical detail to consider is the choice of the interval length, k . The interval should be large enough so that a proof may always be found if the inequality $\sum m_i b_i \geq Q$, or $\sum m_i b_i \leq Q$ holds. An upper bound for the required k can be obtained by considering the minimum and maximum possible values of $\sum m_i b_i$.

Verification Step: The verification process also follows the process reviewed in Section 3. During the process the investor should also consult the prospectus to obtain the authenticity and completeness of the parameters m_i and Q behind the restrictions Limit_i . Once the proof data is verified to be complete and correct, the investor will know that the claimed statements constraining the assets are correct, relative to the assumption that the commitments themselves were not fraudulently created.

Failures and Disputes: If any verification step fails, then the investor knows that a condition of the investment contract has been breached- this should never happen if the fund manager respects the fund composition restrictions. If there is a legitimate reason for the manager to violate a constraint specified in the contract, the manager should not publish a proof-attempt that will fail, but rather address the problem directly. In case of a legal dispute, the commitments can serve as evidence of the claimed portfolio, and as mentioned above, third parties can assist in such a process.

4.3 Discussion

It is clear that the fund manager and investor will need appropriate infrastructure to fully benefit from this mechanism, so it may be most applicable to large institutional investors. A hedge fund which is able to offer this kind of additional assurance would be compensated with ability to attract greater business, and the service might be reflected in the fee that the fund is able to charge.

The scheme increases the accountability of the fund manager, as the investor will have continuous confirmation that the fund has not left the acceptable risk range. The mechanism we describe is certainly stronger than the reputation and *post-facto* legal based approaches in place today. Through the deliberate specification of acceptable risk bounds in the fund prospectus, the mechanism provides strong incentive for the fund manager to manage the portfolio in a manner which is more closely aligned with the investors' risk preferences. Conversely, it discourages investment behavior that concentrates enormous risk on an unlikely scenario, unless the investor agrees to this kind of gamble.

5 Applications

Portfolio risk depends on the evolving allocation among security types, so we now turn our attention to the task of designing meaningful risk constraints within our framework. These constraints take the form of linear combinations of the asset quantities A_i , and include direct limitations on the portfolio composition, as well as constraints based on factor exposures and scenario analysis. Clearly, not all portfolio risks can be specified in advance (or with linear constraints), so our framework leaves open the possibility of revealing additional portfolio risk information not stipulated in the prospectus.

Individual Asset Bounds: These are simple constraints of the form $b_i \leq Q$, which serve to limit the amount invested in a particular single asset A_i . By using this simple constraint for every potential asset, assurance can be obtained that the fund is not placing a significant bet on the performance of a single security.

Asset Class and Sector Allocation: Organizing the list of assets into sectors, a bound on the total investment in a particular sector can be expressed as $\sum m_i b_i \leq Q$, where m_i are non-zero for the assets within the sector, and represent a weighting according to the asset's price at the fund's inception. Sector allocation statements and proofs relative to *updated* asset prices can also be made, but these bounds can not be contractually guaranteed in the same way.

Asset Features, Short Positions: Following the same technique as for sector allocation, the assets can be grouped in any way desired, and an inequality can be constructed bounding the value invested in such a subgroup. An important example of this might be to group the short positions into a group, and bound the amount of asset shorting. This can be accomplished by listing the short positions as distinct assets, or by using constraints of the form $\sum m_i b_i \geq -Q$. Bounding the acceptive complementary short and long positions limits the risks associated with such extreme leveraging, including *liquidity risk*.

Current Minimum Value: An estimation of current value can be communicated by setting the m_i to be the current price, and the statement $\sum m_i b_i \geq -Q$ can be proved for any value of Q less than the actual sum $\sum m_i b_i$. Since such a statement depends on current prices it can not be rigorously guaranteed in the contract, but it may still be a useful piece of information to relate.

Factor exposures: These bounds rely on risk models which assign each asset A_i a factor exposure $e_{i,j}$ to a particular factor F_j . According to such models, the exposure is an estimation of the sensitivity, $d(\text{value})/d(\text{factor})$, to the factor. To use this kind of constraint, the exposures $e_{i,j}$ for factor F_j should be published in the contract. The aggregate sensitivity of the portfolio to F_j is then $\sum e_{i,j} b_i$, which may be positive or negative. A bound $-Q'_j \leq \sum e_{i,j} b_i \leq Q_j$, provides a guarantee that the portfolio is not too sensitive to the factor F_j . For example, such constraints might be used to limit the interest rate risk that the portfolio is allowed to take, or the amount of credit risk.

Scenario analysis: This kind of bound extends the benefit obtained by considering a single risk factor in isolation. First a set of *scenarios* are selected, denoted S_j , which define a set of potential future trajectories of various economic factors. Next, some model must be used to estimate the value $v_{i,j}$ of each asset under each scenario. The prospectus lists the battery of scenarios, and also lists the expected value of each asset under each scenario, and makes reference to the modeling technique used. Finally, an "acceptable risk" is agreed upon by listing the portfolio's minimum future value under each scenario described in the contract. The expected future value of the portfolio under scenario S_j is simply $P_j = \sum v_{i,j} b_i$, so the bound we are interested in takes the form

$$\sum v_{i,j} b_i \geq SV_j. \quad (8)$$

Note that the validity of this approach does not depend on the choice of model: the values $v_{i,j}$ must be published, and the investor must find them reasonable to accept the contract. Of course, the manager can not guarantee future portfolio values, but he can guarantee that he will never take a position which will assume less than the contractual minimum value under any of the listed hypothetical scenarios, however unlikely he feels that the scenario is.

Such scenarios are idealized, discreet, future possibilities, and the actual outcome is unlikely to closely follow an actual scenario listed. Nevertheless, such bounds are very useful since they force the fund to maintain a composition for which it is not expected to lose too much value under an adversarial scenario.

Trading volume: A final type of bound may be useful to detect a certain type of fraud masquerading as “market timing”, where redundant trades are made not to improve the portfolio’s position, but to earn brokerage fees associated with each trade. To allow a bound on the total trading activity within a fund would require a minor tweak: we provide commitments to the amounts of each asset purchased and sold (these are considered separately, and must each be positive). Then bounds on the total amount of sales (purchases) over some period can also be expressed as linear conditions, and the same types of zero knowledge proofs employed.

6 Implementation

To demonstrate the feasibility of our proposal, we implemented a prototype of our scheme using C, and Shoup’s NTL package [12]. For this prototype we generated parameters p and q to be 1024 bits and 160 bits respectively, and used SHA1 as the hash function. With these parameters, each commitment was 1024 bits long, and each k -bit interval proof was $1664k$ bits long. We set the interval proof length, k , to be 30 bits, which is sufficient for the inequalities we would like to prove. This assumes a precision for m_i and b_i of about 15 bits each; increased precision would unlikely significantly add to the risk information conveyed.

Each interval proof with parameter $k = 30$ requires a few seconds to compute, and can be reduced to less than 1 second when optimized on a standard 2005 model PC. Assuming a portfolio with several thousand assets A_i and 1000 constraints Limit_j , the commitments and zero knowledge proofs can be computed in less than twenty minutes, if we assume a cost of 1 second per constraint proof. Of some concern, the proof material does require a substantial amount of space - about 6 megabytes for the parameters [$k=30$, 1000 constraints]. Elliptic curves, or the techniques in [3] may improve efficiency.

The main conclusion we draw for this experiment is that for a reasonably complex portfolio and set of constraints, the computation can be completed in a matter of minutes, and stored at a reasonable cost. This means that it is feasible to generate and publish the proof data at least once per day, for example, after the major US exchanges are closed.

7 Conclusions

This work has introduced, for the first time, the applications of zero knowledge techniques to the release of investment risk material. It is surprising that the relatively simple and well established cryptographic tools of commitment and interval proofs suffice to construct a mechanism to make portfolio composition assertions which can already communicate the most important types of portfolio risks. This follows from the observation that most of the relevant risk assertions (sector allocation, factor exposure, and scenario analysis) are linear in nature.

The premise behind this work is that a verifiable mechanism to communicate risk will increase the trust between an investor and a fund manager, and ultimately create overall economic value. The scheme we describe lies at the crossroads of cryptography, risk management, law, and trust assessment, and is a novel technique to increase accountability of fund managers to investors. The proposed mechanism consists of a contract between the investor and manager, through which the manager agrees to describe the evolving portfolio in a verifiable way. Effectively, the investor will have a new tool to monitor the manager's trades, and to check that the fund characteristics satisfy the risk preferences specified in the contract.

We contend that hedge funds would be more compelling investments, if their performance were not perceived as a magic black-box, often delivering spectacular returns, but occasionally declaring bankruptcy. Indeed, many hedge fund strategies involve taking large positions in oppositely correlated securities, a configuration designed to achieve probable high returns yet only reveal the risks in case of disaster! Despite the fact that the scheme limits the hedge fund manager's choices, he may be motivated to employ our scheme to attract investors who demand real-time risk-exposure information and additional legal assurance.

Acknowledgments

The author would like to thank Dr. Andrew Lo for motivating discussions.

References

1. Securities and exchange commission : Web article, 2003. URL: <http://www.sec.gov/answers/hedge.htm>.
2. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *1st ACM Conference on Computer and Communications Security*, pages 62–73. ACM Press, 1993.
3. F. Boudot. Efficient proofs that a committed number lies in an interval. In Bart Preneel, editor, *Advances in Cryptology - EuroCrypt '00*, pages 431–444, Berlin, 2000. Springer-Verlag. Lecture Notes in Computer Science Volume 1807.
4. E. F. Brickell, D. Chaum, I. B. Damgård, and J. van de Graaf. Gradual and verifiable release of a secret. In Carl Pomerance, editor, *Advances in Cryptology - Crypto '87*, pages 156–166, Berlin, 1987. Springer-Verlag. Lecture Notes in Computer Science Volume 293.

5. J. Campbell, A. Lo, and C. MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press, New Jersey, 1997.
6. R. Cramer, I. Damgård, and B. Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Y.G. Desmedt, editor, *CRYPTO '94*, pages 174–187. Springer-Verlag, 1994. LNCS no. 839.
7. W. Mao. Guaranteed correct sharing of integer factorization with off-line shareholders. In H. Imai and Y. Zheng, editors, *Proceedings of Public Key Cryptography*, pages 60–71. Springer-Verlag, 1998.
8. T. P. Pedersen. A threshold cryptosystem without a trusted party (extended abstract). In Donald W. Davies, editor, *Advances in Cryptology - EuroCrypt '91*, pages 522–526, Berlin, 1991. Springer-Verlag. Lecture Notes in Computer Science Volume 547.
9. D. Pointcheval and J. Stern. Security proofs for signature schemes. In Ueli Maurer, editor, *Advances in Cryptology - EuroCrypt '96*, pages 387–398, Berlin, 1996. Springer-Verlag. Lecture Notes in Computer Science Volume 1070.
10. Claus P. Schnorr. Efficient identification and signatures for smart cards. In *Proceedings on Advances in cryptology*, pages 239–252. Springer-Verlag New York, Inc., 1989.
11. A. Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979.
12. V. Shoup. Ntl: A library for doing number theory, 2003. URL: <http://www.shoup.net/ntl>.

A Cryptography Details

A.1 Schnorr OR Proof Properties

We review the security properties of the Schnorr OR Proof. These are completeness, zero-knowledge and special soundness. The non-interactive version of the proof, also called a sigma protocol [6], is made non-interactive with the Fiat-Shamir transform. Replacing the role of the verifier with a hash function, the non-interactive proofs are proved secure in the Random Oracle Model [2]. There is no known attack on this proof when the random oracle is replaced with a good (one-way, and collision-free) hash function such as SHA1.

Completeness: For any commitment $C_r(0)$, or $C_r(1)$, such a proof can always be efficiently calculated as follows: If $x = 1$ (so $C = g^1 h^r$), let r_1, c_1, u_2 be random (mod q). Let $a_1 = h^{r_1} C^{-c_1} \pmod{p}$, $a_2 = h^{u_2} \pmod{p}$, $c = \text{hash}(a_1, a_2)$, $c_2 = c - c_1$, and $r_2 = u_2 + c_2 r \pmod{q}$. In the case where $x = 0$, (so $C = g^0 h^r$), let r_2, c_2, u_1 be random (mod q), $a_2 = h^{r_2} C / g^{-c_2} \pmod{p}$, $a_1 = h^{u_1} \pmod{p}$, $c = \text{hash}(a_2, a_1)$, $c_1 = c - c_2$, and $r_1 = u_1 + c_1 r \pmod{q}$.

Zero Knowledge: The interactive proof is special honest verifier zero knowledge. For any C, c a simulator which chooses r_1, c_1, r_2, c_2 at random such that $c = c_1 + c_2$, and computes $a_1 = h^{r_1} C^{-c_1}$ and $a_2 = h^{r_2} C / g^{-c_2}$ perfectly simulates

the honest protocol interaction. The non-interactive proof is zero knowledge in the random oracle model.

Special Soundness: This sketch shows that two accepting protocol interactions $(a_1, a_2; c; r_1, r_2, c_1, c_2)$ and $(a_1, a_2; c'; r'_1, r'_2, c'_1, c'_2)$ for a fixed C with different challenges $\{c_1, c_2\} \neq \{c'_1, c'_2\}$ can be used to compute a witness (x, r) for $C = g^x h^r$. Suppose the challenges differ, so either $c_1 \neq c'_1$ or $c_2 \neq c'_2$. In the first case, $h^{(r_1-r'_1)/(c'_1-c_1)} = C$, and in the second, $h^{(r_2-r'_2)/(c'_2-c_2)} = C/g$. Either way a pair (x, r) satisfying $C = g^x h^r$ is found. By the forking lemma [9], the non-interactive proof is thus sound in the random oracle model.