Arbitrage Opportunities Correspond to Probability Inequality Identities

Matthew Doty

February 4, 2026

Abstract

We consider a fixed-odds gambling market over arbitrary logical propositions, where participants trade bets involving conjunctions, disjunctions, and negations. In this setting, we establish a three-way correspondence between the financial feasibility of trading strategies, the validity of universal probability inequalities, and the solutions to bounded Maximum Satisfiability (MaxSAT) problems.

The central result demonstrates that proving a trading strategy constitutes an arbitrage opportunity (i.e., guaranteeing a risk-free profit regardless of the outcome) is equivalent to proving a specific inequality identity holds for all probability functions, and is computationally equivalent to establishing a lower bound on a corresponding MaxSAT instance. Dually, we show that checking the coherence of a strategy (i.e., ensuring it does not guarantee a loss) also corresponds to verifying a probability identity and bounding a MaxSAT problem from above.

Contents

1 Introduction

theory *Arbitrage-Probability-Correspondence*
imports
Probability-Inequality-Completeness.Probability-Inequality-Completeness
HOL.Real
begin

1.1 Motivation

Consider a *fixed-odds gambling market* where participants trade bets on arbitrary logical propositions.

In this setting, every bet pays out exactly \$*1* if the proposition is true and \$*0* otherwise. Unlike traditional prediction markets like *PredictIt* or *Polymarket*, which usually limit trading to mutually exclusive outcomes, we assume a market that allows bets on any combination of logical operators: *AND* ($\sqcap$), *OR* ($\sqcup$), and *NOT* ($\sim$).

To understand the relationship between market liquidity and probability logic, imagine two events:

- A :: *The NASDAQ will go up 1% by Friday*
- B :: *The S&P500 will go up 1% by Friday*

Suppose the market order book contains the following quotes:

- *ASK* for A at \$*0.40* (Someone is selling/offering a bet on A).
- *ASK* for B at \$*0.50* (Someone is selling/offering a bet on B).
- *BID* for $A \sqcap B$ at \$*0.30* (Someone wants to buy a bet on A *AND* B).
- *BID* for $A \sqcup B$ at \$*0.70* (Someone wants to buy a bet on A *OR* B).

An arbitrageur can exploit these prices to guarantee a risk-free profit.

They act as a *market taker* for the *ASK*s (buying A and B) and as a *market maker* for the BIDs (selling A AND B and A OR B).

The initial cash flow is positive:

$Profit = (BID(A \sqcap B) + BID(A \sqcup B)) - (ASK(A) + ASK(B))$ $Profit = (\$0.30 + \$0.70) - (\$0.40 + \$0.50) = \$1.00 - \$0.90 = \$0.10$

Crucially, this profit is safe regardless of the outcome. The arbitrageur holds long positions in A and B, and short positions in $A \sqcap B$ and $A \sqcup B$.

- If both rise (True, True): The arbitrageur wins \$*2* on longs, pays \$*2* on shorts. Net: \$*0* payout.
- If only one rises (True, False): The arbitrageur wins \$*1* on longs, pays \$*1* on short (the *OR* bet). Net: \$*0* payout.
- If neither rises (False, False): The arbitrageur wins \$*0*, pay \$*0*. Net: \$*0* payout.

  The arbitrage exists because the market prices violate the probability identity:

  $Pr(A) + Pr(B) = Pr(A \sqcap B) + Pr(A \sqcup B)$

  The central result of this work generalizes this intuition:

  Every arbitrage opportunity corresponds to a probability inequality identity.

1.2 Overview of Results

The central result of this work is as follows:

Proving a strategy will always yield a profit (if completely matched) in a fixed-odds gambling market over arbitrary logical propositions corresponds to proving an inequality identity in probability logic, and also corresponds to a bounded MaxSAT problem.

Such strategies are referred to as *arbitrage strategies*.

We also consider the *dual* problem of identifying if a trading strategy will never make a profit. Strategies that will never logically yield a profit are called *incoherent*.

1.3 Prior Work

Two results that appear to be related at first glance are *The Fundamental Theorem(s) of Asset Pricing* (FTAP) [6] and the *Dutch Book Theorem* [1, 3, 4, 5]. While the connection to FTAP is purely superficial, the results are close in spirit to the Dutch Book tradition: we study when a collection of fixed-odds commitments can be combined into a strategy that is guaranteed to lose (or, dually, guaranteed to profit), and we treat such strategies as computational objects.

The Fundamental Theorems of Asset Pricing (FTAP) connect a suitable notion of *no-arbitrage* to the existence of a pricing functional (or, in stochastic settings, an equivalent martingale measure) in an idealized, frictionless market. In their classical formulations, the objects being priced are standard financial assets (e.g., securities or commodities) represented by a spot price or a price process, and the market model abstracts away from microstructure: order placement, order matching, bid/ask discreteness, and fixed-odds quoting are not part of the primitive data. By contrast, we work directly with fixed-odds markets for wagers on arbitrary logical propositions, where the microstructure of how orders compose into strategies is central, and we connect "no-arbitage" strategies to the existence of some scenario where the strategy doesn't always lose, which falls out of a certain bounded MaxSAT calculation.

The Dutch Book literature shares more of our vocabulary. Philosophical treatments emphasize *coherence* and the avoidance of a *bad book*: a collection of bets that guarantees a loss. Following Hájek's terminology [2], one may also speak of *good books*. In this development, we adopt finance-oriented language and refer to these objects as (loss-guaranteeing) *arbitrage strategies*, because they are assembled from posted odds and executed mechanically once the relevant orders are matched. We also work with possibility-style representations in the spirit of Lehman, generalized to any instance of a *classical-logic*.

Our main contribution is not a normative thesis that rational agents ought to conform their degrees of belief to probability theory. Instead, we make explicit a three-way correspondence between:

1. checking whether a bounded family of fixed-odds commitments is coherent (i.e., not loss-guaranteeing),
2. feasibility of a bounded MaxSAT instance derived from the same commitments, and
3. certain inequalities that hold for all probability functions over the same set of propositions.

Operationally, we only require the first criterion: there must exist a scenario in which the strategy does not always lose. The MaxSAT formulation supplies a concrete decision procedure, and the coNP-hardness of the resulting feasibility questions explains why coherence checking is not a task one should expect to perform reliably by hand.

We also study the *dual* problem: identifying strategies that are pure arbitrages (guaranteed nonnegative payoff with strictly positive payoff in some outcome). Such strategies are useful not merely as pathologies, but as mechanisms for creating market depth. Intuitively, they can match *BID* interest in one venue with *ASK* interest in another, improving execution for multiple participants. From a microeconomic perspective, this can increase surplus by enabling trades that would otherwise fail to clear.

2 Fixed Odds Markets

notation *Probability-Inequality-Completeness.relative-maximals* (‹$\mathcal{M}$›)

unbundle *no funcset-syntax*

2.1 Orders and Trading Strategies

In this section, we model a *fixed odds market* where each bet pays out \$*0* or \$*1*, and people make and take bets. For simplicity, we consider *BID* and *ASK* limit orders of a single unit (i.e., trades such that if they match, then they are completely cleared). In an ordinary central limit order book, such *BID* and *ASK* orders would have prices in the interval (*0*,*1*), but we do not make use of this assumption in our proofs, as it is not necessary.

record *'p bet-offer* =
  bet :: *'p*
  price :: *real*

A trading strategy is a collection of *BID* and *ASK* orders that are to be matched atomically.

Making a bet is when you *ask* a bet on the market, while *taking* a bet is when you *bid* a bet on the market.

A *market maker* is one who puts up capital and asks bets, while a *market taker* is one who bids bets.

In a trading strategy, the market participant acts as a market maker for the *ASK* orders they are willing make and as a market taker for the *BID* orders they are willing to make.

record *'p strategy* =
  asks :: (*'p bet-offer*) *list*
  bids :: (*'p bet-offer*) *list*

2.2 Possibility Functions

Possibility functions are states of affairs that determine the outcomes of bets. They were first used in Lehman's formulation of the Dutch Book

Theorem [4]. Our approach diverges from Lehman's. Lehman uses linear programming to prove his result. Our formulation is pure probability logic.

We give our definition of a possibility function as follows:

definition (**in** *classical-logic*) *possibility* :: ($'a \Rightarrow bool$) $\Rightarrow$ *bool* **where**
[*simp*]: *possibility* $p \equiv$
$\neg$ (p $\bot$)
$\wedge$ ($\forall$ φ. $\vdash \varphi \longrightarrow p\ \varphi$)
$\wedge$ ($\forall$ $\varphi\ \psi$. p ($\varphi \rightarrow \psi$) $\longrightarrow p\ \varphi \longrightarrow p\ \psi$)
$\wedge$ ($\forall$ φ . $p\ \varphi \vee p$ ($\sim \varphi$))

Our formulation of possibility functions generalizes Lehman's. Lehman restricts his definition to the language of classical propositional logic formulae. We define ours over any arbitrary classical logic satisfying the axioms of the *classical-logic* class.

definition (**in** *classical-logic*) *possibilities* :: ($'a \Rightarrow bool$) *set* **where**
[*simp*]: *possibilities* = {p. *possibility* p}

lemma (**in** *classical-logic*) *possibility-negation*:
assumes *possibility* p
shows p ($\varphi \rightarrow \bot$) = ($\neg$ p φ)
⟨*proof*⟩

lemma (**in** *classical-logic*) *possibilities-logical-closure*:
assumes *possibility* p
and {x. p x} $\Vdash \varphi$
shows p φ
⟨*proof*⟩

The next two lemmas establish that possibility functions are equivalent to maximally consistent sets.

lemma (**in** *classical-logic*) *possibilities-are-MCS*:
assumes *possibility* p
shows *MCS* {x. p x}
⟨*proof*⟩

lemma (**in** *classical-logic*) *MCSs-are-possibilities*:
assumes *MCS* s
shows *possibility* (λ x. $x \in s$)
⟨*proof*⟩

2.3 Payoff Functions

Given a market strategy and a possibility function, we can define the *payoff* of that strategy if all the bet positions in that strategy were matched and settled at the particular state of affairs given by the possibility function.

Recall that in a trading strategy, we act as a market *maker* for ask positions, meaning we payout if the proposition behind the bet we are asking evaluates to *true*.

Payoff is revenue from won bets minus costs of the *BID*s for those bets, plus revenue from sold *ASK* bets minus payouts from bets lost.

definition *payoff* :: $('p \Rightarrow bool) \Rightarrow 'p\ strategy \Rightarrow real$ (π) **where**
[*simp*]: $\pi\ s\ strategy \equiv$
$(\sum\ i \leftarrow bids\ strategy.\ (if\ s\ (bet\ i)\ then\ 1\ else\ 0) - price\ i)$
$+ (\sum\ i \leftarrow asks\ strategy.\ price\ i - (if\ s\ (bet\ i)\ then\ 1\ else\ 0))$

Alternate definitions of the payout function π are to use the notion of *settling* bets given a state of affairs. Settling is just paying out those bets that came true.

definition *settle-bet* :: $('p \Rightarrow bool) \Rightarrow 'p \Rightarrow real$ **where**
$settle\text{-}bet\ s\ \varphi \equiv if\ (s\ \varphi)\ then\ 1\ else\ 0$

lemma *payoff-alt-def1*:
$\pi\ s\ strategy =$
$(\sum\ i \leftarrow bids\ strategy.\ settle\text{-}bet\ s\ (bet\ i) - price\ i)$
$+ (\sum\ i \leftarrow asks\ strategy.\ price\ i - settle\text{-}bet\ s\ (bet\ i))$
⟨*proof*⟩

definition *settle* :: $('p \Rightarrow bool) \Rightarrow 'p\ bet\text{-}offer\ list \Rightarrow real$ **where**
$settle\ s\ bets \equiv \sum\ b \leftarrow bets.\ settle\text{-}bet\ s\ (bet\ b)$

lemma *settle-alt-def*:
$settle\ q\ bets = length\ [\varphi \leftarrow [\ bet\ b\ .\ b \leftarrow bets\ ]\ .\ q\ \varphi]$
⟨*proof*⟩

definition *total-price* :: $('p\ bet\text{-}offer)\ list \Rightarrow real$ **where**
$total\text{-}price\ offers \equiv \sum\ i \leftarrow offers.\ price\ i$

lemma *payoff-alt-def2*:
$\pi\ s\ strategy = settle\ s\ (bids\ strategy)$
$- settle\ s\ (asks\ strategy)$
$+ total\text{-}price\ (asks\ strategy)$
$- total\text{-}price\ (bids\ strategy)$
⟨*proof*⟩

2.4 Revenue Equivalence

When evaluating a payout function, we can essentially convert *BID* orders to *ASK* orders in a strategy, provided we properly account for locked capital when calculating the effective prices for the new *ASK* positions.

definition (**in** *classical-logic*) *negate-bets* ($\text{-}^{\sim}$) **where**
$bets^{\sim} = [b\ (\!|\ bet := \sim (bet\ b)\ |\!).\ b \leftarrow bets]$

lemma (**in** *classical-logic*) *ask-revenue-equivalence*:
assumes *possibility p*
shows $\pi\ p\ (\!|\ asks = asks',\ bids = bids'\ |\!)$
$= -\ settle\ p\ (bids'^{\sim}\ @\ asks')$
$+\ total\text{-}price\ asks'$
$+\ length\ bids'$
$-\ total\text{-}price\ bids'$
⟨*proof*⟩

The dual is also true: when evaluating a payout function, we can, in a sense, treat *ASK* as *BID* positions with proper accounting.

lemma (**in** *classical-logic*) *bid-revenue-equivalence*:
assumes *possibility p*
shows $\pi\ p\ (\!|\ asks = asks',\ bids = bids'\ |\!)$
$=\ settle\ p\ (asks'^{\sim}\ @\ bids')$
$+\ total\text{-}price\ asks'$
$-\ total\text{-}price\ bids'$
$-\ length\ asks'$
⟨*proof*⟩

3 Arbitrage Strategies

3.1 Introduction

In this section, we consider the problem of computing whether a strategy will always yield a profit. Such strategies are referred to as *arbitrage strategies*.

3.2 Minimum Payoff

When computing whether a strategy is suited to arbitrage trading, we need to know the *minimum payoff* of that strategy given every possible scenario.

definition (**in** *consistent-classical-logic*)
minimum-payoff :: *'a strategy* ⇒ *real* (π_{min}) **where**
$\pi_{min}\ b \equiv THE\ x.\ (\exists\ p \in possibilities.\ \pi\ p\ b = x)$
$\wedge\ (\forall\ q \in possibilities.\ x \leq \pi\ q\ b)$

Since our definition of π_{min} relies on a definite descriptor, we need the following theorem to prove it is well-defined.

lemma (**in** *consistent-classical-logic*) *minimum-payoff-existence*:
$\exists!\ x.\ (\exists\ p \in possibilities.\ \pi\ p\ bets = x) \wedge (\forall\ q \in possibilities.\ x \leq \pi\ q\ bets)$
⟨*proof*⟩

3.3 Bounding Minimum Payoffs Below Using MaxSAT

Below, we present our second major theorem: computing a lower bound to a strategy's minimum payoff is equivalent to checking a bounded MaxSAT problem.

A concrete implementation of this algorithm would enable software search for trading strategies that can convert orders from one central limit order book to another.

As in the previous section, we prove our theorem in the general case of an arbitrary k, but in practice users will want to set $k = 0$ to check if their strategy is an arbitrage strategy.

theorem (**in** *consistent-classical-logic*) *arbitrageur-maxsat*:
$((k :: real) \leq \pi_{min}\ (\!|\ asks = asks',\ bids = bids'\ |\!))$
$= (\ \ MaxSAT\ [bet\ b\ .\ b \leftarrow bids'^{\sim}\ @\ asks']$
$\leq total\text{-}price\ asks' + length\ bids' - total\text{-}price\ bids' - k\ )$
(**is** $(k \leq \pi_{min}\ ?bets) = (MaxSAT\ ?props \leq total\text{-}price\ \text{-} + \text{-} - \text{-} - \text{-})$)
⟨*proof*⟩

4 Coherence Checking

4.1 Introduction

In this section, we give an abstract algorithm for traders to use to detect if a strategy they want to employ will *always lose*, i.e., is *incoherent*.

4.2 Maximum Payoff

The key to figuring out if a trading strategy will not always lose is computing the strategy's *maximum payoff*.

Below, we define the maximum payoff using a definite description.

definition (**in** *consistent-classical-logic*)
maximum-payoff :: *'a strategy* ⇒ *real* (π_{max}) **where**
$\pi_{max}\ b \equiv THE\ x.\ (\exists\ p \in possibilities.\ \pi\ p\ b = x)$
$\wedge (\forall\ q \in possibilities.\ \pi\ q\ b \leq x)$

The following lemma establishes that our definition of π_{max} is well-defined.

lemma (**in** *consistent-classical-logic*) *maximum-payoff-existence*:
$\exists!\ x.\ (\exists\ p \in possibilities.\ \pi\ p\ bets = x)$
$\wedge (\forall\ q \in possibilities.\ \pi\ q\ bets \leq x)$
⟨*proof*⟩

4.3 Bounding Maximum Payoffs Above Using MaxSAT

Below, we present our first major theorem: computing an upper bound to a strategy's maximum payoff is equivalent to a bounded MaxSAT problem.

Given a software MaxSAT implementation, a trader can use this equivalence to run a program to check whether the way they arrive at their strategies has a bug.

Note that while the theorem below is formulated using an arbitrary k constant, in practice users will want to check their strategies are safe by using $k = 0$.

theorem (**in** *consistent-classical-logic*) *coherence-maxsat*:
$(\pi_{max}$ ⦇ $asks = asks'$, $bids = bids'$ ⦈ $\leq (k :: real))$
$= (MaxSAT\ [bet\ b\ .\ b \leftarrow asks'^{\sim}\ @\ bids']$
$\leq k - total\text{-}price\ asks' + total\text{-}price\ bids' + length\ asks')$
(**is** $(\pi_{max}\ ?bets \leq k) = (MaxSAT\ ?props \leq \text{-} - total\text{-}price\ \text{-} + \text{-} + \text{-})$)
⟨*proof*⟩

5 Probability Inequality Identity Correspondence

5.1 Introduction

In this section, we prove two forms of the probability inequality identity correspondence theorem.

The two forms relate to π_{min} (i.e., arbitrage strategy determination) and π_{max} (i.e., coherence testing).

In each case, the form follows from the reduction to bounded MaxSAT previously presented, and the reduction of bounded MaxSAT to probability logic, we established in *Probability-Inequality-Completeness.Probability-Inequality-Completeness*.

5.2 Arbitrage Strategies and Minimum Payoff

First, we connect checking if a strategy is an arbitrage strategy and probability identities.

lemma (**in** *consistent-classical-logic*) *arbitrageur-nonstrict-correspondence*:
$(k \leq \pi_{min}$ ⦇ $asks = asks'$, $bids = bids'$ ⦈$)$
$= (\forall\ \mathcal{P} \in probabilities.$
$(\sum b{\leftarrow}asks'.\ \mathcal{P}\ (bet\ b)) + total\text{-}price\ bids' + k$
$\leq (\sum s{\leftarrow}bids'.\ \mathcal{P}\ (bet\ s)) + total\text{-}price\ asks')$
(**is** $?lhs = \text{-}$)
⟨*proof*⟩

lemma (**in** *consistent-classical-logic*) *arbitrageur-strict-correspondence*:
$(k < \pi_{min}$ ⦇ $asks = asks'$, $bids = bids'$ ⦈$)$
$= (\forall\ \mathcal{P} \in probabilities.$
$(\sum b{\leftarrow}asks'.\ \mathcal{P}\ (bet\ b)) + total\text{-}price\ bids' + k$
$< (\sum s{\leftarrow}bids'.\ \mathcal{P}\ (bet\ s)) + total\text{-}price\ asks')$
(**is** $?lhs = ?rhs$)
⟨*proof*⟩

Below is our central result regarding checking if a strategy is an arbitrage strategy:

A strategy is an arbitrage strategy if and only if there is a corresponding identity in probability theory that reflects it.

theorem (**in** *consistent-classical-logic*) *arbitrageur-correspondence*:
$(0 < \pi_{min} \llparenthesis\ asks = asks',\ bids = bids' \rrparenthesis)$
$= (\forall\ \mathcal{P} \in probabilities.$
$(\sum b \leftarrow asks'.\ \mathcal{P}\ (bet\ b)) + total\text{-}price\ bids'$
$< (\sum s \leftarrow bids'.\ \mathcal{P}\ (bet\ s)) + total\text{-}price\ asks')$
⟨*proof*⟩

5.3 Coherence Checking and Maximum Payoff

Finally, we show the connection between coherence checking and probability identities.

lemma (**in** *consistent-classical-logic*) *coherence-nonstrict-correspondence*:
$(\pi_{max} \llparenthesis\ asks = asks',\ bids = bids' \rrparenthesis \leq k)$
$= (\forall\ \mathcal{P} \in probabilities.$
$(\sum b \leftarrow bids'.\ \mathcal{P}\ (bet\ b)) + total\text{-}price\ asks'$
$\leq (\sum s \leftarrow asks'.\ \mathcal{P}\ (bet\ s)) + total\text{-}price\ bids' + k)$
(**is** *?lhs* = -)
⟨*proof*⟩

lemma (**in** *consistent-classical-logic*) *coherence-strict-correspondence*:
$(\pi_{max} \llparenthesis\ asks = asks',\ bids = bids' \rrparenthesis < k)$
$= (\forall\ \mathcal{P} \in probabilities.$
$(\sum b \leftarrow bids'.\ \mathcal{P}\ (bet\ b)) + total\text{-}price\ asks'$
$< (\sum s \leftarrow asks'.\ \mathcal{P}\ (bet\ s)) + total\text{-}price\ bids' + k)$
(**is** *?lhs* = *?rhs*)
⟨*proof*⟩

Below is our central result regarding coherence testing:

A strategy is incoherent if and only if there is a corresponding identity in probability theory that reflects it.

theorem (**in** *consistent-classical-logic*) *coherence-correspondence*:
$(\pi_{max} \llparenthesis\ asks = asks',\ bids = bids' \rrparenthesis < 0)$
$= (\forall\ \mathcal{P} \in probabilities.$
$(\sum b \leftarrow bids'.\ \mathcal{P}\ (bet\ b)) + total\text{-}price\ asks'$
$< (\sum s \leftarrow asks'.\ \mathcal{P}\ (bet\ s)) + total\text{-}price\ bids')$
⟨*proof*⟩

no-notation *Probability-Inequality-Completeness.relative-maximals* (‹$\mathcal{M}$›)

end

References

[1] B. De Finetti. Sui passaggi al limite nel calcolo delle probabilità. 63:1–12.

[2] A. Hájek. Scotching Dutch Books? 19:139–151.

[3] J. G. Kemeny. Fair bets and inductive probabilities. 20(3):263–273.

[4] R. S. Lehman. On Confirmation and Rational Betting. 20(3):251–262.

[5] F. P. Ramsey. Chapter 3. Truth and Probability. In J. B. Braithwaite, editor, *The Foundations of Mathematics and Other Logical Essays*, number 5 in The International Library of Philosophy $ Philosophy of Logic and Mathematics 32, Philosophy of Logic and Mathematics : In 8 Volumes. Routledge.

[6] H. R. Varian. The Arbitrage Principle in Financial Economics. 1(2):55–72.