

The State Risk Premium

A Missing Entry in the Asset-Pricing Typology

Working Paper, July 2026

Rob Levy, Event Intelligence

Preliminary working draft. Comments welcome.

Abstract

Modern asset pricing has a vocabulary for the premia attached to specific forms of exposure: variance, jumps, skew, credit, duration, liquidity, and aggregate equity risk. It has no corresponding entry for the premium attached to bearing pure terminal-state risk.

History, not theory, explains the omission. Arrow-Debreu state-contingent claims have been the mathematical primitive of modern asset pricing for seventy years, but the prices of such claims were rarely observable in lit markets with institutional depth. State exposure was bundled inside options, futures, swaps, credit instruments, parimutuel pools, and bilateral OTC structures. The price of bearing state risk in isolation could not be measured cleanly.

The emergence of bounded-binary state-contingent claims on lit venues changes this. A contract that pays one unit if a specified state occurs and zero otherwise exposes the market price of terminal-state risk directly. This paper defines the State Risk Premium (SRP) as the observable wedge between the cleared price of such a claim and a reference probability for the underlying state. In its cleanest theoretical form, the reference probability is the discounted risk-neutral probability implied by the broader no-arbitrage system. In empirical applications, the reference probability may be physical, model-implied, cross-instrument-implied, or strip-implied, depending on the exercise.

SRP is consistent with no-arbitrage and offers no free lunch. It is the empirically observable price of market incompleteness in the state-pricing dimension. The premium is bounded by the cost of enforcing the arbitrage against related instruments, whether replicated statically through option-implied vertical spreads or carried dynamically as the equivalent gamma and vega exposure in an options book. The paper defines the construct, develops a three-component decomposition, explains its relationship to the variance risk premium and Arrow-Debreu state pricing, and sets out the empirical agenda required to measure it.

Keywords: State risk premium; bounded-binary state-contingent claims; Arrow-Debreu state prices; variance risk premium; prediction markets; event contracts; market microstructure; risk-neutral probability.

JEL Classification: G12, G13, G14, D84.

I. The Open Slot

Asset pricing has accumulated a vocabulary for risk premia over fifty years. Each term names a specific economic phenomenon: a compensation demanded by the marginal participant in a specific market for bearing a specific exposure that the instrument cannot separate from its other exposures.

The Equity Risk Premium (Mehra and Prescott 1985) names the compensation for bearing aggregate market exposure. The Term Premium names the compensation for bearing duration exposure. The Credit Risk Premium names the compensation for bearing default exposure. The Variance Risk Premium (Carr and Wu 2009; Bollerslev, Tauchen, and Zhou 2009) names the compensation for bearing variance exposure over a path. The Jump Risk Premium (Bollerslev, Gibson, and Zhou 2011) names the compensation for bearing discontinuous price movements. The Skew Risk Premium names the compensation for bearing asymmetric tail exposure. The Liquidity Risk Premium names the compensation for bearing the risk of unfavorable execution under stress.

There is no entry in this list for the premium attached to bearing pure terminal-state risk. No one overlooked it; there was nothing to observe. State exposure in conventional instruments was always bundled with something else.

A call option bundles state exposure with variance exposure: the seller bears not only the risk that the underlying finishes above the strike but the realized variance of the path. A futures contract bundles state exposure with carry and roll. A credit default swap bundles state exposure with recovery uncertainty and counterparty risk. A weather derivative bundles state exposure with model risk and settlement opacity. In every conventional instrument, the price of bearing state risk could only be inferred indirectly, by stripping the other exposures away from a bundled price. The stripping was imprecise enough that the residual could never be cleanly observed.

The common feature of these instruments is that each one entangles two distinct bets: which terminal state obtains, and how far the underlying travels to get there. A call holder is long the state (finishing above the strike) and simultaneously long the magnitude of the excursion beyond it. The two cannot be separated inside the instrument, and the price of bearing state risk is therefore never quoted on its own. A fully-collateralized bounded-binary claim performs that separation by construction. Because its payoff is a fixed unit on one side of the threshold and nothing on the other, flat on both sides with all the action at the threshold, the underlying's magnitude drops out of the payoff entirely: the contract pays the same dollar whether the state is realized by a hair or by a mile. What remains is exposure to which state obtains and nothing else. That step structure, not the cap, is what isolates terminal-state risk; the lit venue is what then makes its price observable.

The Arrow-Debreu state pricing framework (Arrow 1953; Debreu 1959) defined the theoretical object: a security paying one unit of account in a specific state of the world and zero in all others. The price of such a security, in a complete market with frictionless trading, equals the risk-neutral probability of that state, discounted by the risk-free rate. Arrow-Debreu securities were the mathematical object that made the theorems of asset pricing work. They were not, until recently, an instrument class whose prices were publicly observable.

Bounded-binary state-contingent claims have existed in various forms for decades. Binary FX options have traded over the counter since at least the 1980s. Cash-or-nothing digitals on equity underlyings have been priced bilaterally between dealers and clients for as long. Parimutuel structures including sports betting markets are economically binary state-contingent claims and have operated at scale for over a century. What changed recently is where they trade: lit venues with continuous order books, publicly observable depth, and timestamped trade data. Polymarket on Polygon, Kalshi under CFTC regulation, and emerging institutional venues list contracts paying \$1 if a specified state occurs and \$0 otherwise, fully collateralized at trade inception, with the marginal participant's cleared price visible to anyone watching the market. The scale is now material rather than experimental. Combined monthly trading volume across Kalshi and Polymarket rose from under \$5 billion in September 2025 to roughly \$24 billion by April 2026 (Pew Research Center 2026, analyzing data from The Block), a level that exceeds the roughly \$14 billion per month wagered through legal sportsbooks in the United States during 2025. Single events concentrate that depth. Polymarket's market on the 2024 United States presidential election recorded roughly \$3.6 billion in cumulative trading volume on the presidential race alone (CNBC 2024), and recurring macro events now clear at meaningful size, with single-meeting federal funds rate contracts on Kalshi reaching volumes near \$100 million for an individual FOMC decision (Diercks, Katz, and Wright 2026). On permissionless venues the on-chain figure combines secondary-market trading with primary share minting and is best read as on-venue notional rather than conventional turnover, but the order of magnitude is no longer in question.

The migration of binary contracts from OTC and parimutuel forms to lit venues is what makes SRP empirically observable at institutional scale. The premium has existed wherever binary contracts have traded; participants in OTC binary FX, in bilateral cash-or-nothing structures, and in parimutuel markets have all faced reservation prices for bearing state exposure that differed from the relevant reference probability. But the pricing, the depth, and the flow were all private, and the premium could not be measured by anyone outside the bilateral relationship or the parimutuel pool operator. The recent migration to lit venues with publicly observable order books is what converts SRP from a phenomenon that exists in the dark into a phenomenon that can be measured, analyzed, and integrated into the broader asset-pricing framework.

There is now a class of publicly priced instruments at meaningful institutional scale whose marginal participants bear terminal-state exposure unbundled from variance, carry, credit, or recovery. The compensation those participants demand is empirically observable in the contract price. The compensation is not VRP, since the marginal participant is not necessarily running a variance-exposed book. Nor is it carry: the contract is fully collateralized. And it is not credit

compensation, because clearing and collateralization leave little counterparty risk to price. What remains is the compensation for bearing the realization of a terminal state. That compensation is the State Risk Premium.

II. Formal Definition

Let O denote a state-contingent event with a binary outcome: state S is either realized or not. Let the market price, $P_{\text{market}}(S)$, denote the price at which a bounded-binary contract paying \$1 if S occurs and \$0 otherwise clears in an open market. Let $\Pi(S)$ denote a reference state probability for S .

The reference state probability is the probability benchmark against which the binary-market price is compared. In its cleanest theoretical form, $\Pi(S)$ is the discounted risk-neutral probability implied by the broader no-arbitrage system:

$$\Pi(S) = e^{-rT} \cdot Q(S)$$

But the empirical exercise does not require every application to begin with a directly observed risk-neutral probability, because $Q(S)$ is not directly observable. Depending on the state being studied, $\Pi(S)$ may be estimated from an option strip, a futures-implied distribution, a related strip of adjacent binary contracts, a structural model, or realized physical frequencies adjusted for known premia. The theoretically preferred reference is risk-neutral. The empirically observable object is often a binary-price wedge versus a reference state probability.

The State Risk Premium is therefore defined as:

$$SRP(S) := P_{\text{market}}(S) - \Pi(S) \quad (1)$$

In the risk-neutral formulation:

$$SRP(S) := P_{\text{market}}(S) - e^{-rT} \cdot Q(S) \quad (2)$$

For practical purposes at short duration, where discounting is negligible:

$$SRP(S) \approx P_{\text{market}}(S) - Q(S)$$

A simple example makes the object concrete. Suppose a one-week SPX cash-or-nothing binary pays \$1 if the index settles above a specified strike at expiry and \$0 otherwise. The contract clears at 18 cents. A reference probability constructed from the relevant option strip, which

prices the same terminal state through the risk-neutral density implied by its strike structure, implies a 17 percent probability after adjustment for the term structure and skew of implied volatility. The observed SRP is therefore +1 probability point. The number is not a forecast error, because the event has not yet resolved; it is the market-clearing wedge between the lit binary price and the chosen reference probability at the time of measurement. If the reference probability is risk-neutral, the +1 point is the theoretically preferred estimate of SRP. If the reference probability is model-implied or physical, the +1 point should be labeled accordingly.

SRP can be positive, zero, or negative. A positive SRP indicates that the binary market clears above the reference state probability. A negative SRP indicates that the binary market clears below the reference state probability. A zero SRP indicates that the market clears at the reference probability, with no observable state-risk wedge under the chosen benchmark.

The definition is intentionally modular. The observed wedge may reflect risk preferences, capital cost, adverse-selection concerns, hedging operations, immediacy provision, or a combination of these effects, which the following sections separate conceptually and propose to measure.

What SRP is / what SRP is not.

SRP is the revealed compensation for bearing terminal-state exposure in a bounded-binary claim. Three things it is not: a violation of no-arbitrage, a free lunch, and prediction-market bias with a new name. What it is: the market-clearing wedge that becomes observable when a lit binary venue reveals the marginal participant's reservation price for bearing a state-contingent payoff.

The definition is parallel in spirit, but not identical in measurement, to the Variance Risk Premium. VRP is usually estimated as the wedge between option-implied risk-neutral expected variance and expected or realized physical variance. SRP can be estimated analogously against physical realized frequencies, but the cleaner theoretical object developed here is narrower: the wedge between a lit binary market's cleared state price and the reference state price implied by the broader no-arbitrage system. The logical structure carries over; the exposure being priced does not.

The separation property. For an unbounded claim with payoff $g(S_T)$, exposure to the terminal state and exposure to the magnitude of the move are borne jointly and cannot be priced independently. For a bounded-binary claim that pays one dollar if S_T exceeds K and zero otherwise, the payoff is invariant to magnitude: the derivative of the payoff with respect to the excursion past K is zero everywhere except at K . The contract therefore strips the underlying's magnitude from the payoff and exposes the marginal participant to terminal-state risk alone. The residual variance of the binary payoff (0.25 at a fifty-cent contract, Section III) is not magnitude exposure to the underlying; it is the intrinsic variance of the isolated state bet itself. SRP is the premium on that isolated object.

III. Why SRP Exists

SRP is neither a behavioral anomaly nor a market inefficiency. The premium follows from a structural feature of how state-contingent claims clear in actual markets. This section develops why the construct exists by examining the marginal participant's reservation price for bearing state exposure.

The marginal participant's reservation price.

The market price of a bounded-binary contract is set by the marginal participant on each side of the order book. The marginal buyer is the participant whose reservation price is the highest bid in the book at any given moment; the marginal seller is the participant whose reservation price is the lowest offer. The market clears where these meet, and the cleared price reflects both reservation prices.

A reservation price for bearing state exposure is not the same as the reference state probability. In the clean theoretical formulation, the reference probability is the discounted risk-neutral probability under a specific equivalent martingale measure. In empirical work, it may be extracted from options, futures, adjacent binary strips, models, or realized frequencies. The reservation price of an actual participant in an actual market reflects the participant's own situation: their utility function, their cost of capital, their hedging operations, their concerns about adverse selection, and the spread they require to compensate them for providing immediacy.

These five factors are the structural sources of SRP. SRP exists because each of them can produce a wedge between the participant's reservation price and the relevant reference state probability, and the wedge persists in equilibrium because the marginal participant in actual markets is not the risk-neutral idealization the theory assumes.

Utility function.

A risk-averse participant facing a binary outcome demands compensation for bearing the variance of the payoff, even on a contract whose terminal payoff has no variance in the conventional path-dependent sense. A 50-cent contract that pays either \$1 or \$0 has a payoff variance of 0.25 regardless of the path the underlying takes between trade and expiry. The variance of the payoff is intrinsic to the binary structure. A risk-averse seller of the contract bears this variance and demands a premium proportional to their risk aversion. A risk-averse buyer also bears variance (they pay 50 cents and may receive \$1 or \$0) and demands a discount.

The risk-aversion premium operates on both sides of the market. In a market with symmetric risk preferences across buyers and sellers, the two effects offset and the contract clears near the risk-neutral probability. In a market with asymmetric risk preferences (more risk-averse sellers than buyers, or vice versa), the contract clears at a price that reflects the imbalance.

The favorite-longshot bias documented in the prediction market accuracy literature (Wolfers and Zitzewitz 2006) is one manifestation of this asymmetry. Longshot contracts (low-probability events) require sellers to bear the variance of a high-magnitude loss with low probability. Risk-averse sellers demand a premium that exceeds the risk-neutral probability of the event. Favorite contracts (high-probability events) produce the symmetric effect on the buy side, with risk-averse buyers demanding a discount.

In closed venues where the bookmaker or operator is structurally positioned as the marginal seller of all longshots and the marginal buyer of all favorites (traditional sportsbooks, OTC binary structures with a single dealer counterparty, parimutuel pools), the two effects compound rather than offset. The operator demands the longshot premium on every longshot ticket and applies the favorite discount on every favorite ticket. The favorite-longshot bias measured in these venues is large and persistent.

In deep lit venues with broad two-sided participant access, the two effects partially neutralize. A market maker who quotes both sides of the order book takes on the longshot-buyer economic position when filling a favorite buyer (because selling the favorite is economically equivalent to buying the longshot) and takes on the favorite-buyer economic position when filling a longshot buyer (because selling the longshot is economically equivalent to buying the favorite). When flow is balanced across favorite buyers and longshot buyers, these two exposures partially offset within the market maker's book. The bias does not disappear (the marginal participant still requires the premium, regardless of which side they are taking) but its magnitude compresses as the venue's depth and two-sided participation grow.

Capital cost.

A participant taking a position in a bounded-binary contract posts collateral equal to their maximum loss. The collateral is locked until the contract resolves. The participant bears the opportunity cost of capital that could otherwise be deployed elsewhere.

The capital cost is non-zero even for participants with deep balance sheets. The marginal use of capital matters: a participant who could earn the risk-free rate on Treasury bills bears at minimum the risk-free rate as the opportunity cost of posted collateral. A participant facing a higher hurdle rate on internal capital allocation bears a correspondingly higher cost. The capital cost enters the reservation price as a discount to bid (the buyer demands the contract cheap enough to compensate them for the locked capital) or a premium to offer (the seller demands the contract expensive enough to compensate them for the locked capital).

For short-dated contracts on liquid underlyings, the capital cost is small (basis-point-scale on a few-day position at the risk-free rate). For longer-dated contracts or contracts on illiquid underlyings with higher hurdle rates, the capital cost can be material. The capital cost component of SRP varies systematically with the contract's duration and the marginal participant's cost of capital.

Hedging operations.

If the marginal participant is hedging a position in the binary contract, they incur the cost of running the hedging program. This includes the cost of executing the hedge, the cost of capital tied up in the hedging instrument, the cost of any imperfection in the hedge (the hedge may not be a perfect offset to the binary exposure), and the operational cost of managing the hedge over time.

The marginal participant in bounded-binary markets today is typically a buy-and-hold-to-expiry participant rather than a dynamic-hedging participant. Their reservation price reflects no hedging operations because they are not hedging. As bounded-binary markets professionalize and attract participants who do hedge their positions, the hedging-operations component of the reservation price becomes more important. The construct does not require the marginal participant to be a dynamic hedger; it accommodates both buy-and-hold and dynamic-hedging operational postures and predicts that the reservation price reflects whichever posture the marginal participant maintains.

Adverse-selection concerns.

The microstructure literature (Glosten and Milgrom 1985; Kyle 1985) establishes that a market maker quoting a continuous bid-ask spread must protect against the risk of being adversely selected by informed traders. The wider the informed-flow share, the wider the spread the market maker requires to break even on the population of orders they encounter.

In bounded-binary state-contingent markets, informed flow can be structurally important. The terminal state of an event is often the kind of thing about which specific participants have genuine information advantages: political analysts on election contracts, central-bank watchers on FOMC contracts, crypto traders on cryptocurrency threshold contracts. Domain knowledge maps cleanly to the binary state in a way it often does not in continuous-price markets where the relevant information concerns timing, magnitude, and path dynamics simultaneously.

The adverse-selection component of the marginal market maker's reservation price reflects the compensation they demand for the risk of being adversely selected by informed traders. It exists in every market with informed flow. Its magnitude in event contract markets is shaped by the structure of the information available about the relevant states.

Immediacy spread.

A market maker provides immediacy: the ability for participants to transact when they wish rather than waiting for a natural counterparty to arrive. The market maker's bid-ask spread compensates them for providing this service, in addition to compensating them for the other costs above. The immediacy spread is the residual market-making compensation that exists even in the absence of risk aversion, capital cost considerations, hedging operations, or adverse-selection concerns. It reflects the operational cost of running the market-making

business and the return required for the market maker to deploy capital to this activity rather than alternatives. Transaction costs belong in this factor and are worth naming explicitly. Venue fees, settlement costs, and the friction of moving collateral are real inputs to the quote: a sophisticated market maker prices them into the bid and the offer the same way they price capital cost, and the taker pays them alongside the spread. The price that actually clears therefore carries the venue's fee structure inside it, on both sides of the book.

The aggregate effect.

These five factors produce a reservation price for the marginal participant that differs from the relevant reference state probability. The market clears at a price that reflects the marginal participant's reservation price. SRP is the difference between the cleared price and the reference state probability. The premium exists because none of the five factors is zero in actual markets, and because their aggregate effect is empirically observable in the wedge between cleared prices and independently estimated probabilities of the underlying states.

The definition is general: it applies to any bounded-binary state-contingent claim with a marginal participant whose reservation price is shaped by the five factors above. The specific magnitude of SRP varies across instruments, venues, and time periods, with the variation determined by the composition of the marginal participants and the conditions under which they form their reservation prices. The decomposition in the next section organizes the five factors into components that can be separately measured and analyzed.

IV. The Decomposition

The aggregate SRP can be decomposed into three structurally distinct components, each organizing a subset of the factors developed in Section III. Each component has its own behavior across market maturity, contract topology, and participant composition.

Component 1: The operational-posture component.

This component reflects the marginal participant's risk preferences, capital cost, and hedging operations as they relate to the participant's overall posture in the market. A participant who is buy-and-hold-to-expiry has different operational requirements than a participant who dynamically hedges. The former bears the variance of the terminal payoff but does not pay realized variance over the path; the latter pays realized variance through their hedging operations but does not bear the same terminal-payoff variance directly.

In current bounded-binary markets, the marginal participant is often buy-and-hold-dominated rather than dynamic-hedging-dominated. The operational-posture component therefore reflects the reservation price of a participant who is bearing terminal-state risk without dynamic-hedging operations. This is typically smaller than the equivalent component in vanilla options markets

(where the marginal participant is dynamically hedging and demands VRP compensation for bearing variance over the path), but it is not zero. The participant still has risk preferences, still incurs capital cost, and may still hedge in other ways.

The operational-posture component shifts as markets professionalize. As bounded-binary markets attract more participants who dynamically hedge, the marginal participant's operational posture shifts toward the dynamic-hedging model, and the component reflects the costs of that posture. The component is also subject to compression through two-sided neutralization in deep lit venues, as discussed in Section III: risk-averse seller and risk-averse buyer effects partially offset each other when broad participant access produces balanced two-way flow, in contrast to closed venues where the operator is structurally on one side of all flow. Both effects (professionalization and two-sided neutralization) operate on the same component but through different mechanisms.

Component 2: The adverse-selection component.

This component reflects compensation for the risk of being adversely selected by informed traders. It exists in every market with informed flow and is not specific to bounded-binary contracts, but its magnitude in event contract markets is shaped by the domain-knowledge mapping between the event and the participant population.

The component has substantial internal structure along two dimensions. The first is cross-sectional variation across event types. Markets referencing events with well-defined information-edge channels (FOMC decisions, earnings releases, geopolitical announcements, regulatory decisions where insider knowledge is plausible) carry larger adverse-selection compensation than markets referencing events with diffuse information (weather, cultural events, sentiment-driven outcomes). The cross-sectional differential reflects the equilibrium informed-flow share, which is itself a function of how much of an information edge a sophisticated participant can realistically construct in the relevant domain. The empirical evidence from recent work on Polymarket and Kalshi (Bürgi, Deng, and Whelan 2025/2026; the longshot spread premium documented in the tick-level Polymarket microstructure literature) is consistent with adverse-selection compensation varying systematically with event category.

The second dimension is time-to-expiry within an event. As an event approaches and as information crystallizes, the adverse-selection regime can intensify dramatically. The hours immediately preceding a scheduled FOMC announcement, after senior Fed officials have made their final public communications, carry a different adverse-selection environment than the same contract two weeks earlier. Conversely, distant horizons typically carry smaller adverse-selection compensation because the information edge has not yet been generated, leaked, or otherwise embedded in the participant population. The time-to-expiry structure of adverse-selection compensation is its own object and is empirically separable from the overall term structure of SRP.

The component does differ systematically across venue types. In closed venues where a single operator is structurally positioned as the counterparty to all flow (sportsbooks, OTC binary structures with a single dealer counterparty), the operator faces informed flow on every trade and must price the full adverse-selection cost into their spreads. In open lit venues with broad participant access, informed flow appears on both sides of the order book, and the per-trade adverse-selection risk borne by any individual market maker is correspondingly smaller. The component therefore exhibits a structural transition between venue types, in addition to its event-type and time-to-expiry heterogeneity within a given venue.

In contrast to the operational-posture component, which compresses systematically as the marginal participant shifts toward a dynamic-hedging dealer who does not need risk-aversion compensation, the adverse-selection component does not compress along the same mechanism. A deeper market with professional participants still faces informed flow, and the professional dealer is still adversely selected by it. The component is not stationary across maturity in a strong sense; it has its own structure. But it does not compress through the same mechanism that compresses the operational-posture component.

Component 3: The structural-hedging-cost component.

This component reflects costs that arise from the contract topology itself, independent of participant population. Bounded-binary contracts have payoff structures that produce specific hedging challenges (concentrated gamma near the strike near expiry, sensitivity to strike discreteness, settlement timing effects). These challenges are properties of the contract, not of the participants who trade it. They contribute to the marginal market maker's reservation price even in a fully professionalized market with sophisticated infrastructure.

The structural-hedging-cost component varies with the contract's specific characteristics (proximity to strike, time-to-expiry, settlement architecture) and with the venue's mechanism design for managing the resulting inventory. The component is present in any bounded-binary market wherever the contract topology generates inventory exposure that is not absorbed by other participants on the book.

The aggregate SRP.

Total SRP is the sum of the three components:

$$SRP = SRP_{\text{posture}} + SRP_{\text{adverse}} + SRP_{\text{structural}} \quad (3)$$

The components have different cross-sectional and time-series signatures. The posture component varies with market maturity and participant population composition and compresses systematically as the marginal participant professionalizes. The adverse-selection component varies cross-sectionally with event-type information characteristics, varies within an event with time-to-expiry, and varies across venue architectures, but does not compress through the same

maturation mechanism that operates on the posture component. The structural-hedging-cost component varies with contract characteristics and venue architecture. The empirical decomposition of an observed SRP into these components is conceptually clean, but the multi-margin identification problem is substantial and the data infrastructure required to identify the components simultaneously is its own research program. The methodology is addressed in Section VI.

V. The Decoupling and the Arbitrage Band

A reader trained in derivatives theory will raise an objection at this point. A bounded-binary contract paying \$1 if $F > K$ at expiry is, by Breeden-Litzenberger (1978), the negative strike derivative of the call price evaluated at K . In a complete market with a dense option strip, the binary is replicable as an infinitesimally tight vertical spread, and no-arbitrage requires the binary price to equal the spread price exactly in the limit $\varepsilon \rightarrow 0$.

If the option strip prices in a VRP, and the binary prices in an SRP, the binary price and the spread price cannot both clear at their respective premia. Arbitrage should drive the two to converge. The State Risk Premium defined above should be identically zero in the limit of complete markets and frictionless trading.

The objection is correct in the idealized limit. What follows explains why it fails in actual markets, and the reasons matter because they define the empirical character of SRP and its relationship to VRP.

Markets are not complete.

The option strip is dense enough to approximate a digital but the approximation is imperfect. Strikes are discrete. Bid-ask spreads on two legs of a vertical spread accumulate. Margin is computed gross rather than net for spread positions. Pin risk affects vertical spreads at expiry in ways that differ in detail from binary pin-risk but are not trivial. The "equivalent vertical spread" is an idealization with $\varepsilon \rightarrow 0$; the actual replication has finite ε and material execution friction. The arbitrage between the binary and the spread is not literally free. It requires putting on two option legs at their own VRP-laden prices and bearing the residual basis. The binary price does not have to equal the spread price exactly; it has to equal it within the cost of the arbitrage.

That cost has a natural expression in greek terms. The replicating position, however constructed, carries the digital's concentrated gamma and vega near the strike, exposures that grow without bound as expiry approaches at the money. The dealer enforcing the band is not paying a spread once; they are warehousing an options-book exposure until resolution, and the width of the band is, in practice, the cost of carrying that greek inventory. This is a number a market maker can mark in the same units used to risk-manage the rest of the book.

Participant populations and books are not fully integrated.

The deeper reason for the decoupling is that the books pricing options and the books pricing bounded-binary contracts price them through hedging operations, capital postings, and adverse-selection exposures that are not identical, even when the firms quoting on either side overlap. The marginal seller of a vanilla option is a professional dealer or vol fund whose business model requires compensation for bearing variance risk because they dynamically hedge the gamma and bear the residual jump risk. The marginal seller of a binary on a lit venue was, until recently, drawn from a meaningfully different population: retail directional traders, crypto-native funds, and information-advantaged participants. That characterization has become less accurate as professional dealers, sometimes the same firms quoting vanilla options on traditional venues, have moved onto the lit binary venues through dedicated liquidity programs. What has not fully converged is the pricing of the two instruments through a single integrated book. Where the binary and the option clear through different clearers, the positions cannot be cross-margined and capital must be posted independently. Where they clear through the same clearer, cross-margin treatment varies by product pair and is not uniformly available. Where they clear in different jurisdictions or settlement systems entirely, for instance a USDC-settled contract on a permissionless venue versus a dollar-margined contract at a US-regulated clearer, capital integration is not available at all. The result is that reservation prices on either side are set with reference to the local hedging constraints, capital costs, and adverse-selection exposure of the book that bears the quote, and these constraints differ in degree across the relevant pairs of books even when they do not differ in kind. The order flow that interacts with each book also still differs in composition. None of these differences is random.

The binary book clears at a price reflecting its marginal reservation price for bearing state risk. The options book clears at a price reflecting its marginal reservation price for bearing variance risk on a path that terminates in the same state. The gap between those two clearing prices is the joint manifestation of SRP and VRP on instruments that price related but not identical exposures. The gap persists because the books, and the flows that interact with each book, do not arbitrage against each other to full integration.

Operational frictions are real.

Polymarket settles in USDC on Polygon. CME options clear through traditional infrastructure with dollar margin. Cross-venue arbitrage requires legal entity structure on both sides, custody of both asset classes, settlement risk between two different clearing systems, regulatory reporting on both sides, and capital posted in both venues simultaneously without cross-margining. The cost of running the arbitrage operationally is substantial and prevents the populations from arbitraging cleanly even when the structural cost gap is observed and understood.

The arbitrage band.

The result is a wide arbitrage band. The binary price floats within the band, set by the marginal binary-participant. The options-replicated spread price sits at the other end of the band, set by the marginal options-participant. SRP and (a transformation of) VRP exist on either side of the band. The band is the price of market incompleteness in the cross-venue, cross-population sense, and the band is wide enough that both premia can be observed simultaneously without violating no-arbitrage.

None of this violates asset-pricing theory; it is the theory operating in its realistic case, where markets are incomplete, populations are segmented, and arbitrage is costly (Shleifer and Vishny 1997). The theory predicts that premia attached to different exposures, traded by different populations, in different venues with material operational frictions, will decouple within a bounded band. Both premia exist, and neither is illusory: they price different exposures, borne by different participants who do not equilibrate against each other.

The band narrows as venues integrate.

The argument above describes the current state of bounded-binary and options markets. The decoupling is structural under current conditions, but the width of the arbitrage band is itself a function of venue integration, and the integration trajectory is in motion. As bounded-binary venues clear through traditional infrastructure, as common clearers cross-margin binary and option positions on related underlyings, as regulatory frameworks treat the two contract types under comparable rules, and as the legal entity structures of dealers consolidate their binary and option books, the operational frictions that maintain the band compress. The Kalshi CFTC framework is one current example: a bounded-binary venue clearing through CFTC infrastructure that some dealers also use for options is structurally closer to the integrated case than a USDC-settled venue on a permissionless chain. The trajectory is not uniform across the binary venue landscape, but the direction is toward integration, and the band narrows in the direction of integration.

The structural prediction follows: as venues integrate, SRP and VRP become more tightly related. In the limit of full integration (common clearer, cross-margin, identical legal and regulatory treatment, dealer books fully unified), the band collapses to the residual width set by the imperfect-replication friction (discrete strikes, two legs of bid-ask, gross margin on the spread, pin risk on the spread legs). In that limit, the binary price and the option-replicated spread price are connected by a tight no-arbitrage relationship, and the joint pricing of SRP and VRP reflects the bundled exposure rather than two independent equilibria. The current observable wedge between binary venue prices and option-strip-replicated prices is therefore a measurable function of venue integration, and the trajectory of the wedge across the integration timeline is a testable empirical object that the framework generates as a prediction.

VI. Measurement

The empirical estimation of SRP requires three things: an observable market price for a bounded-binary contract, a defensible reference state probability for the underlying state, and a methodology for reconciling the two under the relevant no-arbitrage conditions.

The market price is directly observable. Bounded-binary contracts on Polymarket and Kalshi trade with continuous order books and timestamped trade data. The volume-weighted average price over a measurement window, or the mid-price at a specific reference time, provides the empirical market price. One caution applies even at this easy step: the cleared price embeds the venue's transaction costs, because makers price fees into their quotes and takers bear them on execution. An SRP estimate compared across venues with different fee structures should either fee-adjust the observed prices or state explicitly that the fee component is being treated as part of the measured wedge.

The reference state probability is the methodologically difficult input. The cleanest theoretical reference is the discounted risk-neutral probability implied by the broader no-arbitrage system. In practice, four approaches are available, each with strengths and limitations.

Realized frequency under stationarity. For events with sufficient repetition (sports outcomes, repeated economic releases, recurring policy decisions), the long-run realized frequency provides an estimate of the physical probability. This approach is empirically convenient but bundles together the market-clearing SRP with whatever wedge exists between physical and risk-neutral probabilities. It is useful as a benchmark, not as a complete solution.

Cross-instrument consistency. When the same state is referenced by multiple instruments (e.g., an FOMC binary on Kalshi, a Fed Funds futures price, an SR3 option-implied probability), no-arbitrage conditions across instruments allow a reference probability to be extracted from the most liquid and least-premium-contaminated source. The binary-market price is then compared against that reference to yield the observed SRP wedge.

Option-strip consistency. When a vanilla option strip references the same terminal state, the risk-neutral density implied by the strip can be integrated to estimate the reference state probability. This route is natural for financial underlyings, but it requires care because option prices include their own variance, skew, jump, and liquidity premia. The resulting estimate is therefore only as clean as the adjustment from option-implied prices to the desired reference state probability.

Binary-strip consistency. When multiple bounded-binary contracts reference adjacent states (e.g., a strip of SPX cash-or-nothing binaries at adjacent strikes), the cross-sectional shape of the contract prices must satisfy no-arbitrage conditions: probabilities cannot be negative, probabilities across mutually exclusive states must sum to at most one, and the implied density must be non-negative. SRP can be estimated as the deviation between the observed strip and the strip that would clear at the reference distribution implied by a smoothed or interpolated no-arbitrage density.

The four approaches are complementary. A robust empirical estimate would combine them where data permits. The practical point is that SRP is directly observable only after the reference probability is specified; the hard work is anchoring the reference probability, not observing the binary price.

Reference-probability formulations.

The paper's preferred theoretical formulation compares the implied binary price against a discounted risk-neutral reference probability. This choice connects most cleanly to the asset-pricing literature because Arrow-Debreu state prices are risk-neutral state prices in the standard no-arbitrage framework.

An alternative formulation compares the implied binary price against the physical probability of the underlying state. This formulation is empirically convenient because both sides are, in principle, observable: the implied probability is the market price, and the physical probability can be estimated from realized frequencies under stationarity assumptions. But the physical formulation bundles together two distinct effects: the systematic risk-aversion adjustment between physical and risk-neutral probability, and the market-clearing wedge between the reference state probability and the observed binary price.

This paper therefore treats the risk-neutral version as the clean theoretical object and the broader reference-probability version as the practical empirical object. Empirical work should state explicitly which reference probability is being used. If the reference is physical, the measured wedge should be interpreted as SRP plus the relevant risk-premium adjustment between physical and risk-neutral probabilities. If the reference is option-implied, the measured wedge should be interpreted as SRP relative to an option-implied state price that may itself contain variance, skew, jump, and liquidity premia. If the reference is cross-instrument-implied, the measured wedge inherits the assumptions used to clean the reference instrument.

The cross-instrument anchoring problem.

The risk-neutral formulation is theoretically cleaner but introduces a practical difficulty that the physical formulation does not face. The risk-neutral probability Q is not directly observable. In an empirical exercise, the analyst estimates Q from a reference instrument believed to price the same state under a risk-neutral measure that is well characterized. The Fed Funds futures curve provides a reference probability for FOMC outcomes. The vanilla option strip provides a risk-neutral density that can be integrated to yield a probability for a given state. The realized frequency under stationarity provides a physical estimate that must then be adjusted by an assumed risk premium to yield Q .

Each of these reference instruments comes with its own risk-premium-equivalent on its observed price. The Fed Funds futures price can reflect a small but non-zero term premium and policy uncertainty premium. The vanilla option-implied probability reflects option-market premia.

The realized frequency reflects no market premium but requires the stationarity assumption to be defensible. The empirical estimate of SRP on the binary is therefore an estimate of the difference between the binary's implied probability and the reference instrument's implied probability, plus or minus whatever wedge exists between the reference instrument's true risk-neutral measure and its observed price.

Clean cross-instrument anchoring requires either a reference instrument with a well-characterized premium that can be subtracted from its observed price to yield a defensible Q , or a multi-instrument no-arbitrage system that jointly pins down Q across instruments using the no-arbitrage relationships among them. Both routes are workable in principle. Both require careful methodology. The Fed Funds futures route is the cleanest for FOMC-related contracts because the term premium on Fed Funds futures has been studied extensively and can be estimated to reasonable precision. The vanilla option strip route requires a model of VRP to extract Q from the option-implied density, which introduces its own dependence on assumptions about variance dynamics. The realized frequency route requires both stationarity and an estimate of the risk-premium adjustment, the latter of which may be partially the object being measured.

The cross-instrument anchoring problem is the methodologically hardest piece of the empirical estimation of SRP. The construct is well defined and the components are distinguishable in principle, but the practical exercise of separating SRP from the reference instrument's own premium is itself a research methodology that the literature will need to develop. Better to be explicit about that difficulty than to understate it.

Decomposing the components.

The three components of SRP have different empirical signatures and are distinguishable in principle. The operational-posture component varies systematically with market maturity and participant composition; longitudinal comparison across venues at different stages of professionalization isolates it. The adverse-selection component varies along three dimensions developed in Section IV: cross-sectionally with event-type information characteristics, within an event with time-to-expiry, and across venue architectures. The structural-hedging-cost component varies systematically with contract characteristics (proximity to strike, time-to-expiry, settlement architecture) and venue mechanism design.

The richer adverse-selection structure raises the empirical bar for clean identification of the decomposition. Identifying the adverse-selection component alone requires a panel with sufficient cross-sectional variation in event categories, sufficient time resolution within each event, and sufficient venue diversity, while the other components are held constant or jointly estimated. Identifying all three components simultaneously requires a design that supports identification of multiple margins independently, which is a substantially more demanding empirical exercise than a treatment that assumed the adverse-selection component to be roughly stationary across maturity would have implied. The Lauris (2026) pipeline, the Dubach (2026) tick-level Polymarket archive, the Kalshi data behind Bürgi, Deng, and Whelan (2025/2026), and the Diercks, Katz, and Wright (2026) Kalshi work are pieces of the data infrastructure required, but no single dataset currently spans all the margins. A serious empirical

decomposition is a research program that requires assembling a multi-venue panel across these sources or sustained data collection over a window long enough to exploit within-event time-to-expiry variation and across-event category variation simultaneously.

A structural feature of the framework worth flagging.

The operational-posture and adverse-selection components compress through different mechanisms across market maturation, as developed in Section IV. The operational-posture component compresses with professionalization, as the marginal participant shifts toward a dynamic-hedging dealer who does not require the risk-aversion compensation that a buy-and-hold retail participant requires, and through two-sided neutralization in deep lit venues. The adverse-selection component does not compress through the same mechanism: a deeper market with professional participants still has informed flow, and the professional dealer is still adversely selected by it. The component does not compress to zero through maturation, though it may compress to some extent through other mechanisms (such as inventory turnover spreading the adverse-selection cost across a larger volume of flow per market maker).

The implication is structural rather than directly observable in a cross-contract spread comparison. Holding venue maturity, contract topology, and the structural-hedging-cost component constant, the relative weights of the components in observed SRP should shift across maturation toward being adverse-selection-dominated, and the cross-sectional structure of the adverse-selection component (heavier in event categories with strong information-edge channels) should remain visible in mature markets in ways that the operational-posture component does not. The observation does not predict cross-contract spread levels: bid-ask spreads are dominated by capital depth, hedging ecosystem availability, infrastructure quality, and fee structure, none of which the SRP framework directly addresses. An S&P 500 daily-expiring binary will have tighter spreads than a weather contract because it attracts orders of magnitude more capital and has a developed hedging ecosystem, not because its SRP is smaller. The composition-shift observation operates on the decomposition of the SRP wedge, not on the absolute width of bid-ask spreads, and its observation requires the identification methodology developed above to be in place.

The framework's empirical content is therefore conditional on the methodological program rather than freestanding. The decomposition is conceptually clean, the components are in principle distinguishable, and the structural predictions about how they behave across maturation, event type, time-to-expiry, and venue architecture are well-defined. What does not yet exist is the data infrastructure and identification methodology to test those predictions cleanly; building both is its own research agenda.

VII. Relation to the Prediction Market Accuracy Literature

The prediction market accuracy literature has documented systematic deviations between prediction market prices and realized frequencies for two decades, without unifying the deviations under a single risk-premium framework. The construct developed here provides that unification.

The distinction is important. The prediction-market literature asks whether prices are accurate probabilities. The SRP framework asks a different question: what compensation does the marginal participant require to bear a bounded state-contingent payoff? A prediction-market bias is usually evaluated ex post against realized frequencies, while SRP is evaluated ex ante against a specified reference probability. Bias belongs to forecasting; SRP belongs to asset pricing. The two can overlap empirically without being the same object.

The Iowa Electronic Markets work (Berg, Forsythe, Nelson, and Rietz 2008) established that prediction markets aggregate information efficiently in a forecasting sense but documented persistent deviations between prices and realized frequencies that were typically attributed to risk preferences, trader composition, or behavioral biases.

Manski (2006) notes that interpreting prediction market prices as probabilities requires assumptions about risk preferences that are typically violated in practice. He derives bounds on the relationship between prices and probabilities under various preference assumptions but does not name the wedge as a premium.

Wolfers and Zitzewitz (2006) examine the favorite-longshot bias empirically and discuss whether it reflects risk preferences, information asymmetries, or bounded rationality. They find evidence consistent with multiple explanations but do not connect the bias to the broader risk-premium literature in asset pricing.

Snowberg, Wolfers, and Zitzewitz (2013) discuss the relationship between prediction market prices and forecasting accuracy and note that prices can be biased relative to true probabilities for several reasons. They gesture at the gap without naming it.

A related but distinct strand of the literature addresses prediction market design through market scoring rules and parimutuel structures rather than through marginal-participant clearing prices in continuous order books. Hanson (2003, 2007) developed the logarithmic market scoring rule (LMSR), a cost-function-based automated market maker in which a single subsidized market maker quotes prices as a function of accumulated trader holdings. LMSR has been the dominant theoretical and practical mechanism for prediction markets where bootstrapping liquidity at zero volume is the central design constraint, including academic and corporate prediction market deployments and several on-chain implementations. Parimutuel structures, including the long-standing tradition in betting markets and the modern variants in decentralized prediction market mechanisms, clear at pool-ratio prices rather than at marginal-participant clearing prices, with the operator or pool absorbing the price-formation function.

These mechanisms address a different question than the one the SRP framework addresses. LMSR and parimutuel structures are venue-mechanism objects that specify how prices are generated by a single subsidized agent or pool operator. The SRP construct concerns the marginal participant's reservation price in a lit continuous-double-auction venue where prices form through the interaction of many participants. The framework is therefore venue-agnostic with respect to lit CLOB markets but does not transfer cleanly to LMSR or parimutuel mechanisms, where the marginal-participant clearing concept is replaced by a mechanism-determined price function. SRP can be defined and measured at LMSR venues only by reference to the implicit reservation price that traders reveal through their interaction with the cost function, which is a separate methodological exercise. The construct is most directly applicable to the lit CLOB venues that have driven the recent migration of bounded-binary contracts into publicly observable price discovery.

The recent empirical literature on the lit-CLOB venues is consistent with the framework developed here. Bürgi, Deng, and Whelan (2025/2026) document favorite-longshot bias and Kalshi market microstructure evidence, using transaction-level data and the venue's maker-taker identification to study systematic deviations between traded prices and realized win rates in patterns consistent with the operational-posture and adverse-selection components of SRP. Dubach (2026) provides tick-level microstructure evidence from the Polymarket order book, documenting a longshot spread premium, category-conditional effective-spread differences, and depth profiles that map onto features the present framework predicts in the structural-hedging-cost and adverse-selection components. Lauris (2026) develops the broader market-topology characterization on event contract data. Diercks, Katz, and Wright (2026) treat Kalshi contracts as an empirically operational realization of state-contingent claims for macro expectations, constructing probability distributions for macro outcomes directly from the contract strip, which is precisely the framing the present paper develops. The construct does not displace this empirical work; it organizes it.

The State Risk Premium provides one unifying construct that integrates these empirical regularities into the broader asset-pricing framework. Favorite-longshot bias is a manifestation of SRP: a systematic positive premium on longshot contracts and a corresponding pattern on favorite contracts, both arising from the risk-aversion contribution to the operational-posture component. Risk preferences of marginal participants are exactly what the operational-posture component measures. Information asymmetries enter the adverse-selection component. Aggregation efficiency is the question of how SRP shrinks as the market deepens.

The contribution of the present paper is to provide a unified construct that integrates the empirical regularities the literature has documented into the broader asset-pricing framework. The aim is not to replace the existing literature but to gather its findings under a single label that connects them to the broader vocabulary of risk premia in asset pricing.

VIII. Implications

Several implications follow from the SRP framework.

For asset-pricing theory. SRP adds a missing entry to the asset-pricing typology. Every form of exposure that the existing literature has named (variance, jump, skew, equity, term, credit, liquidity) corresponds to a specific bundled exposure inside a conventional instrument. SRP corresponds to the exposure that asset-pricing theory has always assumed as the primitive: the price of a state-contingent claim. That primitive has only recently become publicly observable because the instruments that price it cleanly have only recently begun trading on lit venues at meaningful scale. The Arrow-Debreu state pricing framework, which has been the mathematical scaffolding of modern asset pricing for seventy years, becomes empirically operational when SRP becomes empirically measurable. The probability distribution implied by a dense strip of bounded-binary contracts, after accounting for SRP, is a directly observable object rather than a theoretical idealization inferred solely through pricing models.

For the variance risk premium literature. The decoupling argument in Section V has implications for how VRP should be interpreted. The conventional view treats VRP as compensation for bearing variance risk over a path. The decoupling argument suggests an additional reading: part of what appears as VRP in options markets may reflect the cost of being unable to unbundle state exposure from variance exposure inside the options market. A vol-seller is forced to bear both. The compensation they demand reflects the joint exposure. If the state exposure can be separately offloaded through bounded-binary markets, the variance-specific component of VRP becomes more observable in isolation. This suggests a research agenda, connected to existing work on the drivers of the variance risk premium (Drechsler and Yaron 2011): decomposing VRP into a pure-variance component and a state-bundling component, with the latter theoretically related to SRP under no-arbitrage in the limit. This is the separation property (Section II) read from the options side: the options market cannot perform the state/magnitude separation that the bound performs, so the vol-seller is structurally forced to carry the bundle. SRP and the state-bundling component of VRP are therefore two prices of the same separation, one charged for performing it, one charged for being unable to.

For state-contingent payoffs more broadly. If bounded-binary contracts price SRP separately, and the dense strip of such contracts across adjacent states forms an Arrow-Debreu basis, then any payoff function on the underlying terminal state can be synthesized from the strip without going through a variance-laden options market. The risk-neutral distribution implied by the strip, after accounting for SRP, is itself a directly observable object that can serve as an input to a wide range of downstream constructions.

For institutional hedging. Binary hedgers (pension funds with rate exposure, corporates with merger contingencies, macro funds with policy views) have historically had no choice but to access state exposure through variance-laden instruments. They paid VRP as a forced purchase of unwanted variance compensation, because the instruments available to them could not separate the state bet they wanted from the magnitude bet they did not. The emergence of bounded-binary markets at institutional scale allows them to access state exposure directly,

paying SRP instead of VRP. The structural cost differential, where it exists, is the surplus that flows to the binary hedger from the unbundling.

For regulatory and policy treatment. The distinction between SRP and VRP has implications for how event contracts should be classified in regulatory frameworks. Bounded-binary contracts price a different primitive exposure (terminal-state risk, compensated by SRP) than the options they are often compared to (path-dependent variance and convexity exposure, compensated in part by VRP). In integrated markets, the two instruments may be connected by tight no-arbitrage relationships, and in some cases a binary can be approximated by an option spread. But economic relatedness does not make them the same instrument. Bounded-binary contracts expose and clear terminal-state risk directly. Regulatory frameworks that treat event contracts as a distinct instrument class, such as the CFTC's authorization of Kalshi, are structurally consistent with the framework developed here.

For market microstructure. The framework connects the bounded-binary market microstructure to the broader inventory model literature (Stoll 1978; Ho and Stoll 1981) while identifying features specific to the bounded-binary contract topology. The structural-hedging-cost component of SRP, in particular, picks up the inventory exposures that the contract topology generates and that the venue's mechanism design either absorbs or distributes back into bid-ask spreads. These are concrete phenomena that can be analyzed within the same theoretical apparatus that the microstructure literature has developed for other instruments.

IX. Open Questions

The framework developed here raises several open questions for subsequent research.

The behavior of SRP across market maturation. As bounded-binary markets professionalize and deepen, several distinct compression mechanisms operate on the components of SRP. The operational-posture component is expected to shift toward the dynamic-hedging cost incurred by professional dealers as the marginal participant becomes more sophisticated. The same component is also expected to compress through two-sided neutralization as broad participant access lets risk-averse seller and risk-averse buyer effects partially offset each other, in contrast to closed venues where the operator is structurally on one side of all flow. The adverse-selection component does not compress through the same maturation mechanism; informed flow scales with depth, and the professional dealer is still adversely selected. The structural-hedging-cost component varies with venue architecture, with two-way flow on a deep book enabling internal netting within market maker books. The empirical trajectory of each component across the maturation of current and future bounded-binary venues is testable and would benefit from longitudinal study. The expected pattern is asymmetric compression: each component compresses through a different mechanism with a different dependence on depth, professionalization, and architecture.

The term structure of SRP. Short-dated and long-dated binaries likely command different SRP because the cost of capital tied up in the position differs, the information asymmetries differ, and the hedging horizons of the marginal participants differ. The shape of the SRP term structure has implications for how to construct SRP-adjusted probability term structures and for the pricing of multi-tenor binary products.

The time-to-expiry structure of the adverse-selection component. The adverse-selection component of SRP has internal structure that the overall term structure of SRP does not fully capture. As an event approaches and information crystallizes, the adverse-selection environment can shift dramatically within the same contract. The hours preceding a scheduled FOMC announcement carry a different adverse-selection regime than the same contract a week earlier. The same is true of earnings releases, regulatory decisions, and other scheduled information events. Empirically isolating the time-to-expiry profile of adverse-selection compensation, distinct from the time-to-expiry profile of the other components, is a substantial research agenda requiring event contract data with sufficient time resolution and a panel design that separates this margin from the others. The Kalshi maker-taker data and the Polymarket tick-level archives are starting platforms for this work, but the full identification methodology is itself an open question.

The cross-event-type structure of the adverse-selection component. Adverse-selection compensation varies systematically across event categories with different information-edge channels. Markets referencing events where domain expertise translates into information edge (rate decisions, regulatory actions, geopolitical announcements) carry larger adverse-selection compensation than markets referencing events with diffuse information (weather, cultural events, sentiment-driven outcomes). The empirical mapping between event category and adverse-selection compensation, controlling for the other components, is its own research agenda. The mapping has practical implications for venue product design and for the categorization of binary contracts in regulatory frameworks.

The moneyness structure of SRP. The risk-aversion contribution to the operational-posture component varies with moneyness through the relationship between probability levels and the variance of the binary payoff. The structural-hedging-cost component varies systematically with proximity to the strike near expiry. The shape of the moneyness structure connects to the favorite-longshot bias literature and to the pricing of out-of-the-money binary contracts.

The interaction between SRP and VRP in integrated venues. The decoupling argument in Section V establishes that SRP and VRP coexist in different venues without violating no-arbitrage. The interaction between them within a single venue that lists both binary contracts and options on the same underlying is more subtle. The hypothesis worth testing is that an integrated venue with cross-margining drives SRP and VRP into a tight no-arbitrage relationship, with joint pricing reflecting the bundled exposure. This is testable once such venues exist at scale.

The relationship between SRP on a binary and SRP on the cross-instrument strip. When a state is referenced by multiple contracts (binary, option spread, futures basis), each instrument prices a related but not identical exposure, and the SRP on each may differ. The relationship between instrument-specific SRPs and the underlying state-specific SRP is an open theoretical question with implications for cross-instrument arbitrage and for the construction of state-price estimates from multi-instrument data.

Empirical decomposition methodology. Section VI identifies that the three components of SRP have different empirical signatures and are in principle distinguishable, but the practical methodology for the decomposition requires development. The richer adverse-selection structure articulated in Section IV means the empirical identification problem is multi-margin: the panel must support independent variation in event category, time-to-expiry within event, venue architecture, and contract topology, with the other margins held constant or jointly estimated. Identifying any one component cleanly requires solving for the others. No single existing dataset spans all the required margins; the methodology depends on assembling a multi-venue panel from sources such as the Polymarket tick-level archive, the Kalshi maker-taker data, the Lauris event contract pipeline, and the data infrastructure behind the Diercks, Katz, and Wright (2026) Kalshi work. The empirical decomposition is therefore a research program in its own right, not a methodology that can be executed on existing data with existing tools.

Cross-instrument anchoring methodology. The reference state probability is not directly observable in its clean risk-neutral form and must be estimated from a reference instrument or a multi-instrument no-arbitrage system, as developed in Section VI. Each reference instrument carries its own risk-premium-equivalent on its observed price, which means the empirical estimate of SRP may bundle the construct with whatever wedge exists between the reference instrument's true risk-neutral measure and its observed price. The methodology for cleanly separating SRP from the reference instrument's own premium is the methodologically hardest piece of the empirical program. Promising routes include Fed Funds futures anchors for FOMC-related contracts (where the term premium has been well studied), multi-instrument no-arbitrage systems that jointly identify Q across reference instruments, and stationarity-based estimates of physical probabilities combined with separately estimated risk-premium adjustments. Each route requires methodological development that the literature has not yet produced.

X. Conclusion

This paper defines a construct that the asset-pricing typology has been missing. The State Risk Premium is the empirically observable wedge between the cleared market price of a bounded-binary state-contingent claim and a reference probability for the underlying state. In its cleanest theoretical form, the reference probability is the discounted risk-neutral probability implied by the broader no-arbitrage system. In empirical applications, it may be physical, model-implied, cross-instrument-implied, or strip-implied, depending on the exercise. SRP exists because the marginal participant in the binary market has a reservation price for bearing state exposure that

differs from that reference probability, with the reservation price reflecting the participant's utility function, capital cost, hedging operations, adverse-selection concerns, and immediacy requirements. The construct is parallel in spirit to the Variance Risk Premium but prices a different exposure on a different instrument traded by a different population.

The principal contribution is the construct itself, its formal definition, its three-component decomposition, and its location in the asset-pricing typology. The construct provides a unifying framework for two decades of empirical work in the prediction market accuracy literature. It establishes that the Arrow-Debreu state pricing framework, which has been the mathematical scaffolding of modern asset pricing for seventy years, becomes empirically operational when SRP becomes empirically measurable. It implies a research agenda for decomposing VRP into pure variance and state-bundling components. It has direct implications for the synthesis of state-contingent payoffs from binary strips, for the economics of institutional binary hedging, and for the regulatory treatment of event contracts.

The construct is defined, the decomposition is specified, and the measurement methodology is sketched, with the empirical identification of the components and the cross-instrument anchoring of the risk-neutral probability requiring their own methodological programs that the paper sets in motion. The empirical agenda is open. The relationship to existing premia is established, the decoupling argument explains why the premium persists without violating no-arbitrage, and the connection to the Arrow-Debreu framework places the object in the deepest layer of asset-pricing theory. The construct becomes empirically observable in the specific class of instruments (bounded-binary state-contingent claims with material institutional depth on lit venues) that the financial system has only recently begun to supply at scale. The construct is now available for empirical work, theoretical extension, and integration into the broader asset-pricing framework.

References

- Arrow, K.J. (1953). "Le rôle des valeurs boursières pour la répartition la meilleure des risques." *Économétrie* (Colloques Internationaux du Centre National de la Recherche Scientifique) 11, 41-47. Translated and reprinted as "The Role of Securities in the Optimal Allocation of Risk-Bearing." *Review of Economic Studies* 31(2), 91-96 (1964).
- Berg, J.E., Forsythe, R., Nelson, F., and Rietz, T. (2008). "Results from a Dozen Years of Election Futures Markets Research." In Plott, C. and Smith, V. (eds.), *Handbook of Experimental Economics Results*. Elsevier.
- Bollerslev, T., Gibson, M., and Zhou, H. (2011). "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities." *Journal of Econometrics* 160(1), 235-245.
- Bollerslev, T., Tauchen, G., and Zhou, H. (2009). "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies* 22(11), 4463-4492.
- Breeden, D.T. and Litzenberger, R.H. (1978). "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business* 51(4), 621-651.

Bürgi, C., Deng, W., and Whelan, K. (2025/2026). "Makers and Takers: The Economics of the Kalshi Prediction Market." CESifo Working Paper No. 12122 / University College Dublin Centre for Economic Research Working Paper 2025/19.

Carr, P. and Wu, L. (2009). "Variance Risk Premiums." *Review of Financial Studies* 22(3), 1311-1341.

CNBC (2024). "Polymarket Plans U.S. Return After Signaling Trump Election Win, Founder Says." CNBC, November 7.

Debreu, G. (1959). *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. Yale University Press.

Diercks, A.M., Katz, J.D., and Wright, J.H. (2026). "Kalshi and the Rise of Macro Markets." Finance and Economics Discussion Series 2026-010, Board of Governors of the Federal Reserve System.

Drechsler, I. and Yaron, A. (2011). "What's Vol Got to Do with It." *Review of Financial Studies* 24(1), 1-45.

Dubach, P.D. (2026). "The Anatomy of a Decentralized Prediction Market: Microstructure Evidence from the Polymarket Order Book." arXiv:2604.24366.

Glosten, L.R. and Milgrom, P.R. (1985). "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics* 14(1), 71-100.

Hanson, R. (2003). "Combinatorial Information Market Design." *Information Systems Frontiers* 5(1), 107-119.

Hanson, R. (2007). "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation." *Journal of Prediction Markets* 1(1), 3-15.

Ho, T. and Stoll, H.R. (1981). "Optimal Dealer Pricing under Transactions and Return Uncertainty." *Journal of Financial Economics* 9(1), 47-73.

Kyle, A.S. (1985). "Continuous Auctions and Insider Trading." *Econometrica* 53(6), 1315-1335.

Lauris (2026). "Seeing Like a Market: Event Contracts and Market Topology." Working paper.

Manski, C.F. (2006). "Interpreting the Predictions of Prediction Markets." *Economics Letters* 91(3), 425-429.

Mehra, R. and Prescott, E.C. (1985). "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15(2), 145-161.

Pew Research Center (2026). "Trading Volume on Prediction Markets Has Soared in Recent Months." Short read by J. Gramlich, May 27. Analysis of data from The Block.

Shleifer, A. and Vishny, R.W. (1997). "The Limits of Arbitrage." *Journal of Finance* 52(1), 35-55.

Snowberg, E., Wolfers, J., and Zitzewitz, E. (2013). "Prediction Markets for Economic Forecasting." In Elliott, G. and Timmermann, A. (eds.), *Handbook of Economic Forecasting*, Vol. 2A. Elsevier.

Stoll, H.R. (1978). "The Supply of Dealer Services in Securities Markets." *Journal of Finance* 33(4), 1133-1151.

Wolfer, J. and Zitzewitz, E. (2006). "Interpreting Prediction Market Prices as Probabilities." NBER Working Paper No. 12200.