Markets for Price Risk*

Baiyang Han[†] Junyi Hu[‡] Anthony Lee Zhang[§]

July 2026

Abstract

Many financial contracts – futures, options, and swaps – are written on endogenous prices rather than primitive states of the world. How well can such markets approximate Arrow (1964)'s state-contingent contracts? We develop a tractable equilibrium model in which spot markets allocate goods efficiently but create wealth exposures to equilibrium prices. In our model, futures and variance swaps achieve *price-contingent completeness*: a planner restricted to price-contingent wealth transfers cannot Pareto-improve on equilibrium. Low-dimensional price derivatives therefore implement optimal risk sharing over the component of wealth risk spanned by prices, while leaving residual basis risk uninsurable. Our framework unifies classic results on futures markets and identifies a new role for variance markets in socially beneficial risk sharing.

*We thank David Baqaee, Doug Diamond, Will Diamond, Darrell Duffie, Piotr Dworczak, Arpit Gupta, David Hirshleifer, Alex Imas, Baiyun Jing, Michael Johannes, Ralph Koijen, Dan Luo, Liming Ning, Hrishikesh Relekar, Yijing Ren, Daniel Rock, Milena Wittwer, Xiaobo Yu and seminar participants at Columbia University and UChicago Booth for helpful comments. We are grateful to Jingxi Li for excellent research assistance.

†Stanford University Graduate School of Business; byhan@stanford.edu

‡USC Marshall School of Business; junyihu@marshall.usc.edu

§University of Chicago, Booth School of Business; anthony.zhang@chicagobooth.edu.

1 Introduction

Arrow (1964) modelled financial assets as state-contingent contracts. But oil futures, cattle options, and interest rate swaps are not contracts on exogenous states of the world: their payments depend on the market prices of goods such as oil, cattle and bonds. Financial markets are thus largely low-dimensional markets for price risk. How well do these markets approximate the idealized financial markets studied by Arrow?

To address this question, we build an analytically tractable model of risk sharing in incomplete markets using price derivatives. Spot markets allocate goods efficiently, but generate distortions to risk-sharing. Price derivatives are imperfect financial contracts, which efficiently share the component of wealth risk spanned by spot prices, but leave consumers exposed to unspanned basis risk. We characterize equilibrium trades, risk premia, and welfare in closed form, showing the distinct roles played by futures and variance swap markets. Futures and variance swaps in our model implement a notion we call *price-contingent completeness*: a social planner restricted to implementing price-contingent wealth transfers cannot Pareto-improve on the equilibrium allocation generated by futures and variance swaps. Our framework unifies classic results on futures markets, and introduces a new argument for the role that variance markets play in socially efficient risk sharing.

Our model has two goods: "money", and a real commodity, representing a good such as oil, wheat, or steel. Agents have a concave production technology which converts the commodity into money-equivalents: thus, conditional on shock realization, preferences are quasilinear over the commodity and money. Agents are risk-averse over monetary wealth, with CARA utility with potentially different levels of risk aversion. The only source of uncertainty is agents' random commodity endowments, which we call inventory shocks.

The first-best outcome in our model is characterized by two conditions. First, for any realization of inventory shocks, market outcomes must be *allocatively efficient*: commodities must be allocated to agents to maximize money-equivalent wealth. Second, *risk sharing* should be optimal: shocks to social wealth induced by random aggregate commodity endowments should be shared efficiently across agents, as in Borch (1962) and Wilson (1968). The first-best outcome can be decentralized using Arrow (1964) securities. *Spot markets*, which open after uncertainty is resolved, achieve allocative efficiency; rich *state-contingent contract* markets then implement first-best risk-sharing.

In the absence of financial markets, spot markets optimize allocative efficiency, but do not efficiently share risk between agents across states of the world. We highlight three kinds of distortions to risk-sharing. First, *inventory volatility* exposures are distributed inefficiently across agents: spot "market makers" have *positive* volatility exposure, leaving spot "liquidity takers" with more negative exposure than the economy at large. Second, realizations of inventory shocks represent wealth shocks, which spot markets do not distribute across agents. Third, expected spot trades generate directional price risk – sellers (buyers) have positive (negative) spot price exposures – which spot markets do not efficiently net across agents.

We proceed towards analyzing equilibrium with *price-linked derivative contracts*, developing two main technical results.

In our setting, each consumer's uncertain wealth has a common component driven by endogenous spot good prices, and an idiosyncratic component driven by her own inventory shock. By projecting inventory shocks onto prices, we obtain a "price-contingent certainty equivalent" result: agents behave as if they *only* face price risk, with idiosyncratic risk absorbed into a price-contingent wealth penalty. The original incomplete-markets problem thus reduces to a complete-markets problem, where the spot price is the only source of uncertainty.

Second, classically, optimal financial asset positions must equate "risk-neutral" – that is, marginal-utility-weighted – expected payoffs of all securities to their prices. In our CARA-normal environment, consumers' risk-neutral distributions over spot prices are Gaussian, with means and variances "tilted" relative to the physical spot price distributions. This delivers closed-form expressions for financial asset demands based on moments of the Gaussian distribution. We show that an agent's optimal asset purchases are driven by contract prices and the first- and second-order exposures of her wealth to spot prices; by analogy to option "Greeks", we call these exposures *delta* and *gamma*.

Futures contracts are markets for *delta*, allowing consumers to trade linear wealth exposures to the spot price. Deltas are driven by three components. The first is *expected spot purchases*: consumers who expect to be net sellers (buyers) tend to benefit (suffer) when the spot price is high. The second is *inventory shock variances*: since the spot price is driven by aggregate inventory, a consumer's realized inventory tends to be low when the spot price is high. The third is *basis risk*: idiosyncratic shocks are more costly when the spot price is greater in absolute value. All three forces can contribute to dispersion in consumers' deltas, and thus welfare-improving futures trade. As in classical models, futures let spot buyers and sellers neutralize their offsetting spot price exposures. But futures trade can also arise if certain consumers have highly variable inventory shocks, since long futures positions naturally hedge inventory risks, generalizing insights from the classic futures literature (McKinnon, 1967; Rolfo, 1980).

The futures price is pinned down by the economy's aggregate delta. When the good is valuable, high spot prices coincide with low aggregate inventory, so aggregate wealth is low precisely when the spot price is high. The economy thus has negative aggregate delta: consumers in the aggregate want to be long the futures contract to hedge wealth risk. Market clearing thus implies a negative futures risk premium: the futures price exceeds the expected spot price, so long futures positions make negative expected returns. Interestingly, generalizing Hirshleifer (1990), directional hedging of expected spot trades does not affect the futures risk premium in our model. Spot markets always clear: the economy as a whole cannot buy or sell the spot good on average. Thus, sellers' short hedges always exactly offset buyers' long hedges, creating trade volume but no net demand pressure on futures prices.

A *variance swap* pays consumers based on the squared deviation of the spot price from its mean, allowing consumers to trade *gamma*, second-order exposure to spot prices. We provide a novel intuition for how price variance markets improve social welfare. *Liquidity takers* in spot markets – for example, oil consumers who inelastically demand random quantities of oil – are naturally short volatility: extreme shocks raise their price impact costs and lower their expected wealth. *Market makers* in spot markets – for example, oil producers who supply oil in response to consumers' demand shocks – are naturally *long* volatility: their profits rise when shocks are more extreme. Variance swaps therefore play the same role for volatility risk that futures play for direction risk: they allow liquidity takers and market makers to neutralize their offsetting *variance* exposures, reducing wealth risk and increasing expected utility for both parties.

In our model, the economy is always short gamma: society's aggregate wealth is concave in the spot price. Intuitively, society's marginal utility for the good is downward-sloping: negative inventory shocks harm society more than positive inventory shocks benefit it. Market incompleteness increases this concavity, since higher price variance increases the wealth cost of unhedgeable inventory shocks. The economy thus wants to be long variance in the aggregate to hedge wealth risk. Market clearing thus implies a negative variance risk premium: variance swaps trade above their $\mathbb{P}$-fair value, so long variance positions earn negative returns in equilibrium. In equilibrium, variance markets reallocate society's net gamma exposure to market participants according to their risk aversions.

Futures and variance swaps are optimal price-contingent transfer tools in our model: a social planner who can only implement transfers that are functions of spot market prices cannot Pareto-improve on the equilibrium allocation generated by futures and variance swaps. We term this property *price-contingent completeness.* Technically, agents' $\mathbb{Q}$-measures are Gaussian in our model, and a normal distribution is fully characterized by its mean and variance, implying that first- and second-order price derivatives fully align consumers' entire conditional-certainty-equivalent pricing kernels (marginal utilities), which is exactly the condition for optimal risk sharing in the sense of Borch (1962). Using this result, we analytically characterize the welfare losses from the incompleteness of price derivatives markets relative to first-best risk sharing.

1.1 Related Literature

We contribute to literatures on equilibrium in incomplete markets, futures contracts, and options markets.

General Equilibrium in Incomplete Markets (GEI). We develop a tractable GEI model of standardized contracts written on an endogenous commodity price. Whereas the existing GEI literature generally specifies asset payoffs directly as functions of primitive states, we impose the institutional restriction that financial contracts can condition only on the equilibrium spot price. We find that price-linked contracts efficiently share the price-spanned component of wealth risk, but leave consumers exposed to residual basis risk.

Futures. Suppose a farmer produces a constant x units of crop, and sells at a random price p; clearly, the optimal hedge is to short x futures contracts, which fully neutralizes price risk. Hedging is more complicated when x is random, and may be related to p through market clearing. Characterizing optimal hedging and equilibrium outcomes in futures markets with incomplete risk sharing is a technically difficult problem that has been studied in a sizable literature (Hirshleifer, 1977; Baesel and Grant, 1982; Britto, 1984; Weller and Yano, 1987; Hirshleifer, 1988*a*,*b*, 1989*a*,*b*; Ekeland, Lautier and Villeneuve, 2019). To maintain tractability, existing models typically restrict heterogeneity to a small number of agent types and states; existing models generally do not deliver closed-form expressions for equilibrium outcomes – trade, contract prices, and welfare – under rich heterogeneity. More fundamentally, the literature has paid little attention to a central welfare question: how closely can futures markets approximate the Arrow benchmark of perfect risk-sharing? A separate set of papers studies complete-market models of futures trading (Black, 1976; Breeden, 1980; Hirshleifer, 1990); these approaches deliver predictions about futures pricing, but imply trivial or degenerate equilibrium trading volume and abstract from the incompleteness that makes the Arrow benchmark relevant.

Our contribution to this literature is a fully analytic incomplete-markets model of futures markets. With an arbitrary finite number of heterogeneous agents and continuously distributed inventory shocks, our model delivers closed-form and economically intuitive expressions for trade volume, risk premia, and welfare gains relative to the first-best benchmark. This allows us to unify and generalize many insights from the existing literature, such as the relationship between optimal hedges and demand elasticity discussed in McKinnon (1967) and Rolfo (1980).

More recently, Lyu, Rostek and Yoon (2023) study futures contracts in a dynamic market with strategic traders. Their futures contract is payoff-equivalent to the underlying asset at maturity and is redundant under perfect competition, but becomes nonredundant when imperfect future risk-sharing and cross-round price inference distort spot trade. Futures mitigate this distortion and improve welfare with independent inventory shocks, although the welfare effect can reverse when inference about interdependent values is beneficial. Our mechanism is different: agents in our model are price takers, and futures trade because spot-market activity creates unshared wealth exposures to an endogenous commodity price.

Options. The pricing of options is a cornerstone of modern asset pricing theory (Black and Scholes, 1973) and the subject of an extensive theoretical and applied literature. Yet the literature has paid less attention to a simple equilibrium question: what drives trade in options markets, and what risk-sharing role do such markets serve? The answer to this question is intuitively clear for *futures* contracts: spot goods buyers and sellers have offsetting first-order spot-price exposures, which can be mutually hedged using futures. It is less clear how spot markets would ever give agents natural exposure to price *variance*, to be traded in markets for second-order spot-price risks.

Our paper delivers a simple qualitative answer to this question. Walrasian equilibrium in spot markets makes "market makers" naturally long volatility, and "liquidity takers" naturally short. Consistent with classic narratives regarding the role of risk-sharing markets in society (Bernstein, 1996; Shiller, 2012), variance derivatives enable a *division of labor* between physical production and financial risk-bearing: commodity producers can optimally process inventory in spot markets, and then use variance swaps to offload the resulting financial risks to professional speculators.

Brennan and Cao (1996) and Chabakauri, Yuan and Zachariadis (2022) motivate volatility derivative trading for financial assets through *asymmetric information*: in rational expectations equilibrium, private information gives informed and uninformed investors different valuations of quadratic payoffs, creating an information-driven motive to trade volatility exposure. Our mechanism is complementary but distinct: we abstract from informational frictions and show that spot-market trade itself creates heterogeneous gamma exposures that agents share through variance swaps. Information-based mechanisms may be especially relevant for options on financial assets, whereas our mechanism may be more relevant for commodity options and other derivatives on spot good prices.

Gârleanu, Pedersen and Poteshman (2009) analyze a model of demand-based option pricing in which risk-averse dealers are compensated in equilibrium for absorbing exogenous end-user option demand. Our focus is different: we analyze a microfounded model of how spot markets endow agents with natural exposures to spot price variance, which can then be traded in derivative markets. In our setting, equilibrium contract trade and risk premia are pinned down by economic primitives – inventory risks, spot demand slopes, and risk aversions – rather than by an exogenous vector of option demands.

We also relate more loosely to a broader literature on variance swaps. Martin (2011) introduces simple variance swaps and the associated SVIX index, and Martin (2017) uses SVIX to infer expected market returns. Our focus is different: we study the trade, pricing, and welfare effects of quadratic claims on endogenous spot prices in incomplete markets.

Price-Linked Contracts. A small literature studies contracts indexed by future equilibrium prices. Svensson (1981) and Heal (2017) take agents' price beliefs as primitives, while Kurz and Wu (1996) impose Rational Beliefs – distinct from rational expectations – in a one-good stock-price model with complete price-state claims. Henrotte (1996) constructs an endogenous no-arbitrage price support and proves temporary-equilibrium existence. Correia-da-Silva (2015) shows that price-contingent deliveries reproduce Arrow-Debreu allocations when prices separate primitive states, but does not characterize optimal risk sharing through price-contingent transfers when prices pool states.

When earlier authors obtain welfare conclusions, they rely on some combination of complete contracting over price states, state-separating prices, and subjective beliefs that need not match the equilibrium price law. We instead impose deliberately stylized assumptions to derive a common, model-consistent price law. We show that equilibrium with price-linked contracts exhausts all gains available through price-contingent transfers; however, basis risk generally prevents first-best risk sharing. We call this qualified welfare property *price-contingent completeness.*

Financial Product and Market Design. A related literature shows that the payoff span alone does not determine whether financial innovation affects equilibrium or welfare. Rostek and Yoon (2021) study span-preserving changes in market-clearing protocols; Rostek and Yoon (2024) compare innovation in market-clearing technology with the introduction of synthetic products; and Rostek and Yoon (2025) characterize efficient product design when securities clear independently. Limited cross-security demand conditioning changes price inference even with competitive investors, while large traders add an endogenous price-impact channel. Products that are redundant in payoff terms can therefore be nonredundant as market institutions.

In contrast, our model features a single nonstorable spot good and otherwise standard competitive markets, shutting down both the price-impact and restricted-demand-conditioning channels. Instead, derivatives are simply nonredundant in payoff terms: contracts written on the future spot price are the only way agents can trade exposure to second-period inventory and price risks.

1.2 Paper Summary

The paper proceeds as follows. Section 2 introduces the model. Section 3 characterizes first-best outcomes. Section 4 analyzes outcomes in spot markets, together with complete Arrow securities, and then in the absence of financial markets. Section 5 develops the "conditional certainty equivalent" methodology we use to analyze derivative markets. We analyze futures contracts in Section 6, and variance swaps in Section 7. We discuss our model assumptions and relationship to other literature in Section 8, and conclude in Section 9. We relegate some longer and less informative proofs and derivations to Appendix A.

2 Model

Notationally, we will use bold symbols for vectors, for example writing $\boldsymbol{x}$ to mean the vector $(x_1, \ldots, x_N)$.

There are N "types" of consumers indexed by i, with a representative consumer of each type who behaves competitively, ignoring price impact.[1] For expositional simplicity, we will refer to the representative consumer of type i as simply "consumer i". Consumers have CARA utility over monetary wealth, with possibly different risk aversions α_i:

$$U_i\left(W_i\right) = -e^{-\alpha_i W_i} \tag{1}$$

There are two goods: money and a single commodity. Consumer i is endowed with an initial constant amount m_i of money, and all consumers can hold arbitrarily large positive or negative positions in goods and money. Each consumer i has a quadratic "production technology", which converts any positive or negative quantity y_i of goods into wealth:

$$W_i = m_i + \underbrace{\psi y_i - \frac{y_i^2}{2\kappa_i}}_{Production\ Technology} \tag{2}$$

[1]This is equivalent to assuming there is a unit measure of identical atomistic consumers of each type, who behave competitively because their trades are too small to move prices.

CARA utility implies that money endowments m_i have no effect on i's behavior, since m_i simply scales $U_i(W_i)$ in (1) by a constant factor; thus, we proceed to set $m_i = 0$ for all i, so we can write W_i simply as a function of y_i:

$$W_i(y_i) = \psi y_i - \frac{y_i^2}{2\kappa_i} \tag{3}$$

Wealth $W_i(y_i)$ consists of a linear component ψy_i, which pays the consumer ψ per unit of the commodity, and a quadratic "inventory cost" component $\frac{y_i^2}{2\kappa_i}$, which implies that the marginal monetary value of the good is decreasing in the amount of the good held. Consumers with higher κ_i have lower inventory costs, and thus more elastic demand for the good. We will allow the edge case of $\kappa_i = 0$: we interpret such a consumer as having no capacity to hold the commodity, so she has perfectly inelastic demand for exactly $y_i = 0$ units of the commodity, and attains $-\infty$ wealth under any other value of y_i.

We call $W_i(y_i)$ a "production technology" because it is intuitive to think of y_i being literally transformed into units of consumable wealth. After "transformation" of y_i, the economy reduces to a single-good problem: each consumer has some amount of produced wealth, which can be redistributed across consumers arbitrarily, since money is tradable and consumers have deep pockets. Of course, it is isomorphic to think of $W_i(y_i)$ as a preference function for y_i rather than a production technology; in these terms, consumer i gets utility equivalent to having $W_i(y_i)$ extra dollars from having y_i units of the commodity.

Uncertainty in the model arises from *inventory shocks*: consumer i begins with a random endowment x_i of the commodity. We assume the x_i are independent normal random variables, with means μ_i and variances σ_i^2 that may vary across consumers. If i receives inventory shock x_i, and purchases q_i of the commodity at price p per unit, her final wealth is:

$$W_i = \psi(x_i + q_i) - \frac{(x_i + q_i)^2}{2\kappa_i} - pq_i \tag{4}$$

We will impose the normalization that the mean of the aggregate inventory shock is 0:

$$E\left[\sum_{i=1}^{N} x_i\right] = \sum_{i=1}^{N} \mu_i = 0 \tag{5}$$

Appendix A.1 shows that (5) is purely a normalization and has no substantive content, because any nonzero average in the aggregate inventory shock can be absorbed into the definition of ψ.

There are two periods. The first period is a market for *risk*: consumers may trade *financial securities* which alter their endowments of goods or money in future states of the world. In the following sections, we will analyze three financial market structures: complete financial markets with Arrow securities; no financial markets; and commodity price-linked derivative contracts, which we will show constitute an incomplete financial market. We ignore consumption in the first period, so all asset trades in the first period transfer consumption across future states of the world.

Before the second period begins, consumers' inventory shocks x_i are realized. The second period is a market for *goods*, or in traditional terms, a *spot market*: conditional on financial security trades in period 1 and inventory shock realizations $x_1 \dots x_N$, consumers trade money for the commodity. We assume the commodity is nonstorable, ruling out the possibility of intertemporal self-insurance through inventory accumulation, in order to isolate the role of financial contracts in sharing risk.

3 The First-Best Outcome

Conditional on any vector of inventory shocks $\boldsymbol{x}$, the social planner can freely reallocate commodities across consumers; that is, the social planner chooses functions $y_1(\boldsymbol{x})$ to $y_N(\boldsymbol{x})$, satisfying, pointwise in $\boldsymbol{x}$, the aggregate resource constraint:

$$\sum_{i=1}^{N} y_i(\boldsymbol{x}) = \sum_{i=1}^{N} x_i \tag{6}$$

It is of course equivalent to assume that the social planner chooses net trade amounts $q_i(\boldsymbol{x})$ rather than final inventories $y_i(\boldsymbol{x})$. The social planner can also freely reallocate wealth across agents, pointwise in $\boldsymbol{x}$. Conditional on the planner's choice of final inventories, society's aggregate wealth is:

$$W(\boldsymbol{y}(\boldsymbol{x})) \equiv \sum_{i=1}^{N} W_i(y_i(\boldsymbol{x})) = \sum_{i=1}^{N} \left[\psi y_i(\boldsymbol{x}) - \frac{(y_i(\boldsymbol{x}))^2}{2\kappa_i}\right] \tag{7}$$

The social planner thus chooses final monetary wealths of agents, which we will call $G_i(\boldsymbol{x})$, subject to the constraint, pointwise in $\boldsymbol{x}$, that:

$$\sum_{i=1}^{N} G_i(\boldsymbol{x}) \leq W(\boldsymbol{y}(\boldsymbol{x})) \tag{8}$$

Thus, in sum, the social planner chooses commodity allocations $y_i(\boldsymbol{x})$ and money allocations $G_i(\boldsymbol{x})$, satisfying (6) and (8). An allocation is *Pareto efficient* if it is not expected-utility dominated by some other allocation; formally, under our assumption of CARA utility, $\tilde{G}_i(\boldsymbol{x})$ Pareto-dominates $G_i(\boldsymbol{x})$ if:

$$E\left[-e^{-\alpha_i \tilde{G}_i(\boldsymbol{x})}\right] \geq E\left[-e^{-\alpha_i G_i(\boldsymbol{x})}\right] \quad \forall i, \quad \text{and } E\left[-e^{-\alpha_i \tilde{G}_i(\boldsymbol{x})}\right] > E\left[-e^{-\alpha_i G_i(\boldsymbol{x})}\right] \quad \text{for some } i \tag{9}$$

Notice that, while commodity allocations $y_i(\boldsymbol{x})$ do not explicitly enter into (9), they matter because they constrain money allocations $G_i(\boldsymbol{x})$ through the wealth constraint (8).

Proposition 1. *Pareto-efficient commodity allocations $y_i^*(\boldsymbol{x})$ and money allocations $G_i^*(\boldsymbol{x})$ are characterized by two conditions:* spot market allocative efficiency*, and* optimal risk-sharing*. Spot market allocative efficiency requires that commodity allocations $y_i^*(\boldsymbol{x})$ satisfy, a.s. in $\boldsymbol{x}$:*

$$y_i^*(\boldsymbol{x}) = \frac{\kappa_i}{\sum_{j=1}^N \kappa_j} \sum_{j=1}^N x_j \tag{10}$$

In any efficient spot market allocation, society's aggregate wealth is:

$$W^*(\boldsymbol{x}) \equiv W(\boldsymbol{y}^*(\boldsymbol{x})) = \psi \sum_{i=1}^N x_i - \frac{\left(\sum_{i=1}^N x_i\right)^2}{2\sum_{i=1}^N \kappa_i} \tag{11}$$

Optimal risk-sharing requires that wealth is shared, a.s. in $\boldsymbol{x}$, as:

$$G_i^*(\boldsymbol{x}) = C_i + \frac{\alpha_i^{-1}}{\sum_{j=1}^N \alpha_j^{-1}} W^*(\boldsymbol{x}), \tag{12}$$

where $\sum_{i=1}^N C_i = 0$.

Proof. See Appendix A.3. □

Intuitively, in any Pareto-efficient allocation, spot market commodity allocations $y_i^*(\boldsymbol{x})$ must be efficient: commodities must be distributed to optimally produce money, given consumers' heterogeneous production technologies $W_i(y_i)$. If this were not the case for any $\boldsymbol{x}$, society could simply reallocate goods to generate more wealth, and redistribute this wealth to increase all consumers' utilities in state $\boldsymbol{x}$. Expression (10) states that, since all consumers have quadratic inventory costs, the aggregate endowment $\sum_{j=1}^N x_j$ is simply divided among consumers proportional to their inventory capacities κ_i.

Through the optimal spot market allocations, society simply transforms commodities $\boldsymbol{x}$ into some total monetary wealth $W^*(\boldsymbol{x})$, characterized by (11). Intuitively, when spot markets function optimally, the N consumers' wealth is equivalent to a single representative consumer with inventory capacity:

$$K \equiv \sum_{i=1}^{N} \kappa_i$$

Conditional on optimal spot market allocations, society then faces a simple one-good risk-sharing problem: there is some random total monetary wealth $W^*(\boldsymbol{x})$ which is to be divided amongst risk-averse consumers. Then, Pareto efficiency requires marginal-utility ratios to be equalized across states. Under our assumption of CARA utility, the classic results of Borch (1962) and Wilson (1968) imply that Pareto-efficient allocations redistribute risks in wealth, driven by uncertainty in $\boldsymbol{x}$, affinely according to consumers' risk tolerances, as in (12).

Pareto efficiency is thus a very restrictive criterion in our model: spot market allocations are fully pinned down, and wealths are pinned down across states up to consumer-specific constants. With slight abuse of terminology, we will thus refer to the outcomes described in Proposition 1 as "the first-best outcome" in singular form, ignoring the constant terms in (12).

4 Spot Market Equilibrium

We next solve for equilibrium in spot markets, and illustrate why spot markets fail to implement the first-best outcome.

After inventory shocks x_i are realized, there is no remaining uncertainty, and our setting reduces to a simple quasilinear-utility Walrasian equilibrium, where the only goods are money and the commodity. From (4), i's wealth is:

$$W_i = C_i + \psi (x_i + q_i) - \frac{(x_i + q_i)^2}{2\kappa_i} - pq_i \tag{13}$$

where q_i is the amount of the commodity i purchases at market price p, and C_i is any monetary endowment i may have obtained through financial-asset trading in the first period. Differentiating (13) with respect to q_i, we have:

$$\frac{\partial W_i}{\partial q_i} = \psi - \frac{x_i + q_i}{\kappa_i} - p \tag{14}$$

Expression (14) depends on x_i and q_i, but not C_i: preferences are quasilinear, so there are no income effects. Thus, arbitrary wealth transfers in financial markets have no effect on spot market demand and thus on equilibrium prices and quantities.

W_i is concave in q_i; thus, setting (14) to zero and solving for q_i, we obtain consumers' demand for the good as a function of the spot price p:

$$q_i(p) = -x_i - \kappa_i (p - \psi) \tag{15}$$

Hence, the inventory shock x_i determines the intercept of the demand curve, and inventory capacity κ_i determines the slope.

Spot market equilibrium is characterized by a market-clearing scalar price p. Summing consumers' demand and setting to zero, we require:

$$\sum_{i=1}^{N} [x_i + \kappa_i (p - \psi)] = 0$$

The spot market clearing price $p^{Spot}(\boldsymbol{x})$ is thus simply a function of consumers' inventory shocks:

$$p^{Spot}(\boldsymbol{x}) - \psi = -\frac{\sum_{i=1}^{N} x_i}{\sum_{i=1}^{N} \kappa_i} \tag{16}$$

Intuitively, the equilibrium price deviation from ψ is negative the aggregate inventory shock $\sum_{i=1}^{N} x_i$, divided by the slope of aggregate demand, $\sum_{i=1}^{N} \kappa_i$. Substituting (16) into consumer demand (15) yields consumers' equilibrium inventories:

$$x_i + q_i^{Spot}(\boldsymbol{x}) = \frac{\kappa_i}{\sum_{j=1}^{N} \kappa_j} \sum_{j=1}^{N} x_j \tag{17}$$

That is, i ends up holding a fraction $\frac{\kappa_i}{\sum_{j=1}^{N} \kappa_j}$ of the aggregate inventory shock $\sum_{j=1}^{N} x_j$, implementing the first-best outcome (10). Intuitively, conditional on the realization of inventory shocks, the two-good money-and-commodity market is trivially complete, and the welfare theorems hold. Spot market competitive equilibria are thus *allocatively efficient*: they always allocate commodities in a way which maximizes society's aggregate monetary wealth.

Note also that our normalization in (5) implies that:

$$\sum_{i=1}^{N} E[x_i] = \sum_{i=1}^{N} \mu_i = 0$$

This conveniently implies from (16) that the mean of the spot price is equal to ψ, and also that i's expected spot market purchase quantity is equal to $-\mu_i$, since:

$$E\left[q_i^{Spot}(\boldsymbol{x})\right] = -E[x_i] + \frac{\kappa_i}{\sum_{j=1}^{N} \kappa_j} \sum_{j=1}^{N} E[x_j] = -\mu_i \tag{18}$$

We can calculate the wealth distribution induced by competitive equilibria in spot markets, by plugging equilibrium prices (16) and quantities (17) into consumers' wealth, (13). In Appendix A.2, we show that this simplifies to:

$$W_i^{Spot} = \psi x_i + \frac{\kappa_i}{\sum_{j=1}^N \kappa_j} \frac{\left(\sum_{j=1}^N x_j\right)^2}{2\sum_{j=1}^N \kappa_j} - x_i \frac{\sum_{j=1}^N x_j}{\sum_{j=1}^N \kappa_j} \tag{19}$$

where we have omitted the constant money term, C_i, for convenience. Alternatively, substituting (16) for $p^{Spot}(\boldsymbol{x})$ into (19) and rearranging, we have:

$$W_i^{Spot} = p^{Spot}(\boldsymbol{x})\, x_i + \frac{\kappa_i}{2}\left(p^{Spot}(\boldsymbol{x}) - \psi\right)^2 \tag{20}$$

where for convenience we do not explicitly write the dependence of W_i^{Spot} on $\boldsymbol{x}$.

4.1 Arrow-Debreu Securities and the First-Best Outcome

In the first period, suppose agents can trade Arrow securities which fully span the state space. These securities are contingent on $\boldsymbol{x}$: let $\theta_i(\boldsymbol{x})$ denote the net payoff to i's portfolio when state $\boldsymbol{x}$ is realized. Since there is no first-period consumption, consumers finance securities which pay in some states by selling securities which pay in other states, subject to the usual Arrow-security budget constraints. Because Arrow securities are in zero net supply, their net payoffs must sum to zero across consumers in every state. We will use this complete-market benchmark to characterize the state-contingent transfers required to move from spot-market wealth to first-best risk sharing.

Proposition 2. *The equilibrium with Arrow securities is Pareto-efficient. In particular, equilibrium wealth is a first-best allocation $G_i^*(\boldsymbol{x})$, with the constants in (12) selected by agents' Arrow-security budget constraints. Thus, agent i's Arrow-security demand is:*

$$\theta_i(\boldsymbol{x}) = G_i^*(\boldsymbol{x}) - W_i^{Spot}(\boldsymbol{x}) \tag{21}$$

Proof. Arrow securities span the state space, so markets are complete and the first welfare theorem implies that the equilibrium allocation is Pareto-efficient. Agent i's wealth after Arrow-security trade is $W_i^{Spot}(\boldsymbol{x}) + \theta_i(\boldsymbol{x})$; hence it equals the first-best wealth allocation $G_i^*(\boldsymbol{x})$ selected by the equilibrium. Rearranging gives (21). □

Intuitively, spot and financial markets bring us from autarky to first-best efficiency in two stages. Spot markets implement allocative efficiency, (10); in each state of the world

$\boldsymbol{x}$, spot markets reallocate goods across consumers to optimally convert goods to wealth. This reduces our setting to a one-good problem, where each consumer is endowed with $W_i^{Spot}(\boldsymbol{x})$ dollars in state $\boldsymbol{x}$. Financial markets then implement risk-sharing, (12), optimally sharing state-dependent shocks to aggregate wealth $W^*(\boldsymbol{x})$ across consumers. Equilibrium Arrow-security demands are straightforward: in each state $\boldsymbol{x}$, i buys the difference between her first-best wealth $G_i^*(\boldsymbol{x})$ and her spot-equilibrium wealth $W_i^{Spot}(\boldsymbol{x})$.

Spot market wealths $W_i^{Spot}(\boldsymbol{x})$ generally differ from first-best wealths $G_i^*(\boldsymbol{x})$. One simple way to see this is that consumers' risk aversions α_i influence $G_i^*(\boldsymbol{x})$, but do not enter $W_i^{Spot}(\boldsymbol{x})$. Nothing in spot market clearing therefore ensures that the wealth risks generated by spot trade are shared according to consumers' risk tolerances. We illustrate these differences in two core examples, which we return to when analyzing futures and variance swaps in later sections.

4.2 Market Makers and Liquidity Takers

Suppose there are two types of agents, a "liquidity taker" (T) and a "market maker" (M). T has risk aversion α_T, receives an inventory shock $x_T \sim N(0, \sigma_T^2)$, and has no capacity to absorb inventory, $\kappa_T = 0$. We might think of T as a group of oil consumers, who have perfectly inelastic demand for a noisy quantity $-x_T$ of oil. M has risk aversion α_M, receives no inventory shock, $x_M = 0$, and has positive capacity $\kappa_M > 0$. We might think of M as a group of oil suppliers, who face no uncertainty, but have an upward-sloping production cost curve, allowing them to supply increasing amounts of oil as prices increase.

In equilibrium, the market maker simply buys the entire inventory shock from the liquidity taker:

$$q_T = -x_T,\ q_M = x_T$$

From (12), efficient wealths are:

$$G_T^* = C_T + \frac{\alpha_T^{-1}}{\alpha_T^{-1} + \alpha_M^{-1}} \left(\psi x_T - \frac{1}{2} \frac{x_T^2}{\kappa_M} \right) \tag{22}$$

$$G_M^* = C_M + \frac{\alpha_M^{-1}}{\alpha_T^{-1} + \alpha_M^{-1}} \left(\psi x_T - \frac{1}{2} \frac{x_T^2}{\kappa_M} \right) \tag{23}$$

From (19), equilibrium wealths are:

$$W_T^{Spot} = \psi x_T - \frac{x_T^2}{\kappa_M},\ W_M^{Spot} = \frac{1}{2} \frac{x_T^2}{\kappa_M} \tag{24}$$

Equilibrium wealths feature two distortions.

Wealth Shock Sharing. In our model, the commodity can have some average value – ignoring the concave term, a unit of the commodity is marginally worth ψ dollars – so inventory shocks serve as wealth shocks. When $\psi > 0$, society is wealthier when the endowment x_T is higher; in the first-best outcome, this aggregate wealth shock is split between T and M according to their inverse risk aversions, as indicated by the coefficients $\frac{\alpha_T^{-1}}{\alpha_T^{-1}+\alpha_M^{-1}}$ and $\frac{\alpha_M^{-1}}{\alpha_T^{-1}+\alpha_M^{-1}}$ on x_T in (22) and (23). In contrast, in spot market equilibrium (24), T keeps the entire linear term ψx_T, leaving M no linear exposure to the asset. Somewhat surprisingly, in Section 6, we will show that futures trading can allow M and T to better share aggregate endowment risk.

Volatility Exposures. In the social optimum, all agents have concave exposure to aggregate inventory shocks: the coefficients on x_T^2 in (22) and (23) are respectively:

$$-\frac{1}{2\kappa_M}\frac{\alpha_T^{-1}}{\alpha_T^{-1}+\alpha_M^{-1}}, -\frac{1}{2\kappa_M}\frac{\alpha_M^{-1}}{\alpha_T^{-1}+\alpha_M^{-1}}$$

which add up to society's total exposure $-\frac{1}{2\kappa_M}$. Intuitively, the first-best requires agents to receive constant shares of society's money-equivalent wealth, which is a concave function of aggregate inventory x_T. Interestingly, in spot market equilibrium, the sign of M's volatility exposure is wrong: M actually *benefits* from larger values of x_T^2. Conversely, T has negative volatility exposure of $\frac{1}{\kappa_M}$, which is greater than society's total exposure $\frac{1}{2\kappa_M}$. Liquidity takers essentially pay for volatility both due to the concavity of society's wealth, and also from a "liquidity tax" paid to market makers.

Realized volatility x_T^2 is a random variable that causes W_T^{Spot} and W_M^{Spot} to be risky. The offsetting volatility exposures in (24) thus seem clearly inefficient: if M could pay $\frac{x_T^2}{2\kappa_M} - E\left[\frac{x_T^2}{2\kappa_M}\right]$ to T, T's wealth risk would decrease and M's risk would be eliminated entirely, improving expected utility for both. We will show how this can be accomplished with variance swaps in Section 7.

Physical Capacity, Risk Aversion, and the "Division of Labor". In our model, an agent's demand slope κ_i for the spot good is distinct from their risk aversion α_i over wealth. We view this as realistic, and in fact crucial to the value of financial markets. High κ_i agents may be commodity producers: entities like oil producers or crop farmers, who have made large infrastructure investments to produce or store physical goods, but may be fairly risk-averse over monetary wealth. Low α_i agents may be professional speculators: entities like hedge funds, who have high financial risk tolerance, and are willing to harvest risk premia in various financial markets, but may have no capacity to produce and trade physical commodities.

Authors such as Bernstein (1996) and Shiller (2012) have argued that financial markets enable a *division of labor* between these two groups of firms, allowing commodity producers to specialize in physical processing, offloading the resultant wealth risks to professional speculators. This division of labor is clearly visible in the social planner's outcome, and in idealized financial markets: high-κ_i agents physically bear a large share of aggregate inventory fluctuations, but the resultant financial risks are shifted onto low-α_i agents. In later sections, we will illustrate how it is accomplished in more realistic markets for price risk.

4.3 Buyers and Sellers

Next, we add to T, M two additional consumers, called "buyer" (B) and "seller" (S). B and S have nonrandom inventories, $x_B = -1$, and $x_S = 1$, and have no inventory capacity, $\kappa_B = \kappa_S = 0$. Intuitively, B and S respectively want to inelastically buy and sell a unit of the commodity, and their net demands cancel out.

In spot market equilibrium, goods allocations are:

$$q_T = -x_T,\ q_M = x_T,\ q_B = 1,\ q_S = -1$$

That is, T sells her entire endowment x_T to M, and S sells a unit to B. This allocation generates the same societal wealth as the last example, so in the first-best, this wealth is divided between agents according to their risk aversions, for example:

$$G_S^* = C_S + \frac{\alpha_S^{-1}}{\alpha_T^{-1} + \alpha_M^{-1} + \alpha_B^{-1} + \alpha_S^{-1}} \left(\psi x_T - \frac{1}{2} \frac{x_T^2}{\kappa_M} \right) \tag{25}$$

and likewise for B, T, M. Prices are exactly the same as in the previous example, so the equilibrium wealths of M and T are unchanged from (24). Equilibrium wealths for the buyer and seller are now:

$$W_B^{Spot} = -\psi + \frac{x_T}{\kappa_M},\ W_S^{Spot} = \psi - \frac{x_T}{\kappa_M} \tag{26}$$

Expression (26) illustrates two distortions. First, neither S nor B has any exposure to inventory risk ψx_T and volatility risk $\frac{x_T^2}{2\kappa_M}$. Second, in spot market equilibrium, S inelastically sells a unit of the commodity to B, at a price $\psi - \frac{x_T}{\kappa_M}$ determined by T's inventory shock and M's inventory capacity. Commodity price randomness thus creates variance in W_B^{Spot}, W_S^{Spot}, lowering both parties' ex-ante utility. This is clearly inefficient: these wealth risks are exactly offsetting, so S and B should commit to making a net transfer $\frac{x_T}{\kappa_M}$ to neutralize price risks; equivalently, S and B could simply contract in advance to trade at a fixed price. The classical role of futures contracts is to allow S and B to offset these price risks; we will show in Section 6 how this is accomplished in our model.

5 Conditional Certainty Equivalents and Risk-Neutral Pricing

Building on our analysis of spot markets in isolation, we proceed towards the main goal of the paper: the analysis of realistically incomplete financial markets, consisting of financial derivatives linked to the spot good's price. This section will develop two useful technical results towards this end.

We have thus far worked with inventory shocks x_i; henceforth, we will instead use the demeaned representation:

$$x_i \equiv \mu_i + \epsilon_i$$

where μ_i are constants and ϵ_i are independent mean-zero normal random variables with standard deviations σ_i. From (16), the spot price is then:

$$p^{Spot}(\boldsymbol{x}) - \psi = -\frac{\sum_{i=1}^{N}(\mu_i + \epsilon_i)}{\sum_{i=1}^{N}\kappa_i} = -\frac{\sum_{i=1}^{N}\epsilon_i}{\sum_{i=1}^{N}\kappa_i} \tag{27}$$

where we used the normalization in (5) that $\sum_{i=1}^{N}\mu_i = 0$, that is, the mean inventory shock sums to 0 across agents.[2] Thus, $p^{Spot}(\boldsymbol{x}) - \psi$ is normally distributed, with mean 0 and variance:

$$\sigma_p^2 \equiv \frac{\sum_{i=1}^{N}\sigma_i^2}{\left(\sum_{i=1}^{N}\kappa_i\right)^2} \tag{28}$$

We will also renormalize prices, defining:

$$z \equiv p^{Spot}(\boldsymbol{x}) - \psi \tag{29}$$

where we suppress z's dependence on the vector of inventory shocks $\boldsymbol{x}$ (equivalently, $\boldsymbol{\epsilon}$) for notational convenience. With slight abuse of terminology, we will refer to z as the "spot price" or "normalized spot price". z and ϵ_i are convenient because $(z, \epsilon_1 \ldots \epsilon_N)$ is a vector of mean-0 jointly normal random variables. Substituting z into (20), we can write i's spot equilibrium wealth as:

$$W_i^{Spot} = (\psi + z)(\mu_i + \epsilon_i) + \frac{\kappa_i}{2}z^2 \tag{30}$$

[2]Recall that $\sum_{i=1}^{N}\mu_i = 0$ has no substantive content: Appendix A.1 shows that any nonzero average in μ_i can simply be absorbed into the definition of ψ.

5.1 z-Certainty Equivalents

Consumer i's spot wealth in expression (30) depends on both the common spot price z, which determines the payoffs of spot-price-linked derivatives such as futures and variance swaps; and the idiosyncratic inventory shock ϵ_i. In our model, it turns out we can simply "integrate out" the idiosyncratic ϵ_i terms from utility, transforming the full $(z, \epsilon_1 \dots \epsilon_N)$ problem into the much simpler problem of risk sharing conditional on a single source of uncertainty, the spot price z, which is common to all agents.

From (27), (z, ϵ_i) is jointly normally distributed for any i. We can thus linearly "project" ϵ_i onto z, as follows.

Definition 1. We can write:

$$\epsilon_i = \beta_i z + \eta_i \tag{31}$$

where $\beta_i z$ is the conditional expectation $E[\epsilon_i \mid z]$, and η_i is uncorrelated with z. By computing covariances and variances of z and ϵ_i through (27), we show in Appendix A.4 that:

$$\beta_i = -\left(\sum_{j=1}^{N} \kappa_j\right) \frac{\sigma_i^2}{\sum_{j=1}^{N} \sigma_j^2}, \ \ \nu_i \equiv Var(\epsilon_i \mid z) = \sigma_i^2 \left(1 - \frac{\sigma_i^2}{\sum_{j=1}^{N} \sigma_j^2}\right) \tag{32}$$

Intuitively, the inventory shock ϵ_i is correlated with the price z, because a high inventory shock for any individual consumer predicts slightly lower spot prices, through the spot market clearing condition (27). (31) simply separates the inventory shock into a z-predictable component, $\beta_i z$, and a "fully idiosyncratic" orthogonal residual η_i. β_i is larger in magnitude for agents with comparatively high inventory shock variances, σ_i^2, since the price is driven by these agents' inventory shocks to a greater degree; correspondingly, the residual variance fraction $\frac{\nu_i}{\sigma_i^2}$ is smaller for these agents, since prices "explain" a larger share of their inventory shocks.

A useful identity implied by (32), which we will use in later calculations, is:

$$\sum_{i=1}^{N} \beta_i = -\left(\sum_{j=1}^{N} \kappa_j\right)\left(\sum_{i=1}^{N} \frac{\sigma_i^2}{\sum_{j=1}^{N} \sigma_j^2}\right) = -\sum_{i=1}^{N} \kappa_i \tag{33}$$

Intuitively, when we project the aggregate inventory shock $\sum_{i=1}^{N} \epsilon_i$ onto spot prices, we simply recover the slope of aggregate market demand, $-\sum_{i=1}^{N} \kappa_i$.

Substituting (31) into (30) and rearranging, we can express W_i^{Spot} in terms of a z-spanned component, and an idiosyncratic residual:

$$W_i^{Spot} = \underbrace{(\psi + z)(\mu_i + \beta_i z) + \frac{\kappa_i}{2} z^2}_{Spanned\ by\ z} + \underbrace{(\psi + z)\eta_i}_{Orthogonal\ to\ z} \tag{34}$$

Now, we evaluate the expected utility of an agent i who receives spot wealth W_i^{Spot} plus some arbitrary wealth transfer $t(z)$ which is a function of z (equivalently, of the spot price p^{Spot}).

Lemma 1. *If consumer i receives $\boldsymbol{\epsilon}$-measurable random wealth $W_i^{Spot} + t(z)$, her expected utility is:*

$$\int_{-\infty}^{\infty} -\exp\left[-\alpha_i \left[(\psi + z)(\mu_i + \beta_i z) + \frac{\kappa_i}{2} z^2 + t(z) - \frac{\alpha_i}{2} \nu_i (\psi + z)^2\right]\right] dF(z) \tag{35}$$

Since (35) has no dependence on η_i, i can be thought of as a CARA-utility consumer with wealth that is a z-measurable random variable:

$$(\psi + z)(\mu_i + \beta_i z) + \frac{\kappa_i}{2} z^2 - \frac{\alpha_i}{2} \nu_i (\psi + z)^2 + t(z) \tag{36}$$

Proof. Intuitively, conditional on z, the only randomness in the consumer's utility function is $(\psi + z)\eta_i$; the classical mean-variance property of the CARA-normal system means that the consumer's certainty-equivalent penalty for this randomness is its variance $\nu_i (\psi + z)^2$ multiplied by half i's risk aversion, $\frac{\alpha_i}{2}$.

Formally, applying the law of iterated expectations, a consumer with wealth $W_i^{Spot} + t(z)$ has expected utility:

$$E\left[U_i\left(W_i^{Spot} + t(z)\right)\right] = E\left[E\left[U_i\left(W_i^{Spot} + t(z)\right) \mid z\right]\right]$$

This expands into:

$$= \int \int -\exp\left(-\alpha_i \left[(\psi + z)(\mu_i + \beta_i z) + \frac{\kappa_i}{2} z^2 + t(z) + (\psi + z)\eta_i\right]\right) dF(\eta_i \mid z)\, dF(z)$$

$$= \int -\exp\left(-\alpha_i \left[(\psi + z)(\mu_i + \beta_i z) + \frac{\kappa_i}{2} z^2 + t(z)\right]\right) \left[\int \exp\left(-\alpha_i (\psi + z)\eta_i\right) dF(\eta_i \mid z)\right] dF(z)$$

From (31) and (32) in Definition 1, the conditional distribution $F(\eta_i \mid z)$ is normal with mean 0 and variance ν_i, implying:

$$\int \exp\left(-\alpha_i (\psi + z)\eta_i\right) dF(\eta_i \mid z) = \exp\left(\frac{\alpha_i^2}{2} \nu_i (\psi + z)^2\right) \tag{37}$$

Hence, i's expected utility is equal to (35). (36) is then immediate from inspection. $\square$

Lemma 1 "projects out" the idiosyncratic ϵ_i terms onto the common z term. In doing so, we show that an agent who receives W_i^{Spot}, which is affected by both common price risk z and idiosyncratic risk η_i, is equivalent to a consumer who *only* faces common price risk z, with the modified wealth in (36). The basis risk terms, η_i, collapse into z-contingent "certainty-equivalent" penalties $\frac{\alpha_i}{2}\nu_i(\psi+z)^2$: intuitively, this is just the variance of the residual $(\psi+z)\eta_i$ term in (34), multiplied by the classic $\frac{\alpha_i}{2}$ risk aversion adjustment factor in CARA-normal models. This representation holds conditional on arbitrary z-contingent transfers, whether they are from a social planner, or from financial asset trading.

We will proceed to analyze financial market trading behavior using the representation in (36), which we will call consumers' *z-certainty-equivalent* wealths.

Definition 2. We define i's *z-certainty-equivalent wealth* $\bar{W}_i$, or *z-CE wealth* in short, as (36), ignoring the $t(z)$ term:

$$\bar{W}_i = A_i + \Delta_i z + \Gamma_i z^2 \tag{38}$$

where:

$$A_i = \psi\mu_i - \frac{\alpha_i\nu_i\psi^2}{2} \tag{39}$$

$$\Delta_i = \mu_i + \psi(\beta_i - \alpha_i\nu_i) \tag{40}$$

$$\Gamma_i = \frac{\kappa_i}{2} + \beta_i - \frac{\alpha_i\nu_i}{2} \tag{41}$$

$\bar{W}_i$ is a z-measurable random variable.

By analogy to option pricing practice, Δ_i and Γ_i are "Greeks": the linear ("delta") and quadratic ("gamma") exposures of z-CE wealth $\bar{W}_i$ to the spot price deviation z. Lemma 1 and Definition 2, in eliminating idiosyncratic risk, effectively convert the incomplete-market (z, η_i) problem into a complete-market problem with modified wealths, where z is the only source of uncertainty.

Conceptually, z-CEs are possible because of the level-invariance property of CARA utility. Under other utility functions, certainty equivalents would display *income effects*: $t(z)$ would influence the money-equivalent cost of risk by changing consumers' risk aversion. The additional assumption that η_i is normally distributed leads to particularly simple CE expressions; z-CEs exist, though they are analytically more complicated, for CARA utility and generally distributed noise.[3]

[3] In general, (37) can be written as:

$$E\left[e^{-\alpha_i(\psi+z)\eta_i} \mid z\right]$$

5.2 Risk-Neutral Measures and Derivative Pricing

Using z-CEs, we can take a "risk-neutral distribution" approach to solving for derivative demand. Suppose, in the first period, an agent can commit to buying h_i units of a derivative contract; for each unit purchased, i pays a constant price C in all states of the world, and receives a z-contingent payoff $f(z)$ in return. Since i's derivative payoff is z-contingent, Lemma 1 applies, and i's effective wealth is her z-CE plus her derivative payoffs:

$$\bar{W}_i + h_i (f(z) - C) = A_i + \Delta_i z + \Gamma_i z^2 + h_i (f(z) - C) \tag{42}$$

i's expected utility is thus:

$$E_{\mathbb{P}} \left[U \left(\bar{W}_i + h_i (f(z) - C) \right) \right] \tag{43}$$

where $\mathbb{P}$ denotes the physical probability measure over the spot price z, which has mean 0 and variance (28). Differentiating (43), i's optimal choice of contract position h_i must satisfy the first-order condition:

$$E_{\mathbb{P}} \left[U' \left(\bar{W}_i + h_i (f(z) - C) \right) (f(z) - C) \right] = 0 \tag{44}$$

Rearranging, we have:

$$\frac{E_{\mathbb{P}} \left[U' \left(\bar{W}_i + h_i (f(z) - C) \right) f(z) \right]}{E_{\mathbb{P}} \left[U' \left(\bar{W}_i + h_i (f(z) - C) \right) \right]} = C \tag{45}$$

Under appropriate regularity conditions, the LHS of (45) can be interpreted as a h_i-dependent *change of measure*, which uses $U' \left(\bar{W}_i + h_i (f(z) - C) \right)$ as the pricing kernel, or Radon-Nikodym derivative:

$$\frac{d\mathbb{Q}_i^{h_i}}{d\mathbb{P}} = \frac{U' \left(\bar{W}_i + h_i (f(z) - C) \right)}{E_{\mathbb{P}} \left[U' \left(\bar{W}_i + h_i (f(z) - C) \right) \right]} \tag{46}$$

This can be re-absorbed into the CARA utility function, determining a wealth penalty depending on the *cumulant generating function* of η_i:

$$\frac{1}{\alpha_i} \ln E \left[e^{-\alpha_i (\psi + z) \eta_i} \mid z \right]$$

The normal distribution is the only distribution with a quadratic CGF, so the CE penalty in (37) is only second-order in z when η_i is Gaussian, leading to our "markets for Greeks" representation of derivatives trading. General η_i distributions would introduce higher-order terms in z into certainty equivalents, which agents would require higher-order price derivatives to hedge. While this is a potentially interesting direction in which to extend the model, we leave further analysis to future work.

where we write $\mathbb{Q}_i^{h_i}$ to emphasize that the measure is i-specific, and depends on i's contract purchase quantity h_i. This leads to a simple and classical characterization of i's contract demand.

Lemma 2. *At i's optimal contract position h_i, we must have:*

$$E_{\mathbb{Q}_i^{h_i}}\left[f\left(z\right)\right] = C \tag{47}$$

with $\mathbb{Q}_i^{h_i}$ defined through (46).

Intuitively, consumer i chooses her derivative position h_i, which determines her state-contingent payoffs (42), thus her utility (43) and marginal utility (44), and thus her risk-neutral measure $\mathbb{Q}_i^{h_i}$ through the Radon-Nikodym derivative (46). i adjusts h_i until the classical risk-neutral pricing equation (47) is satisfied: in words, until i's personal risk-neutral probability measure sets the $\mathbb{Q}_i^{h_i}$-expectation of the derivative contract's payoff $f\left(z\right)$ equal to its price C. We will proceed in the following two sections to use this representation to solve for equilibrium in futures and variance swap markets.

6 Futures Markets

We now analyze *commodity futures contracts*. We will analyze a single *cash-settled futures contract* on the commodity. Cash-settled and "physical delivery" contracts are outcome-equivalent in our model; we focus on cash settlement because contracts which transfer wealth instead of goods across states of the world are more intuitive in our setting.[4]

A consumer holding a long (short) contract position promises to pay (receive) some fixed p^c in the second period, and receives (pays) the uncertain spot price p; equivalently, the consumer pays some fixed $z^c \equiv p^c - \psi$ and is paid the normalized spot price z. We thus call

[4]Kyle (2007) shows that cash-settled and physical-delivery derivatives are exactly equivalent whenever a set of "microstructure fungibility" conditions hold; our model satisfies these conditions. Essentially, since spot market outcomes are modelled simply through competitive equilibria, it is equivalent to consumer i's budget set whether she receives some k units of the commodity through a physical-delivery contract, or kp dollars through a cash-settled contract. Trade quantities vary – if the consumer receives k units of the commodity through a futures contract, she must adjust her net trade quantity in spot markets by k – but market clearing in zero-net-supply futures markets implies that these net quantity adjustments cancel out across consumers, leaving net wealths and spot market prices unchanged. This is shown formally in Section 5.4 of Zhang (2022), in a model related to ours. The expositional advantage of using cash-settled contracts in our setting is that, since they only transfer wealth across states of the world, they do not influence consumers' spot market bids, so spot market outcomes are unchanged from the previous section. With physical-delivery contracts, spot wealth outcomes would be identical state-by-state, but consumers' goods trading patterns would be more notationally cumbersome, since consumers would have to adjust spot market bids to account for futures contract "deliveries".

z^c the *contract price*. Let c_i be i's net contract position, and let $c_i(z^c)$ denote i's contract demand as a function of z^c. Contract markets clear through z^c adjusting until aggregate demand for long and short contract positions is equal:

$$\sum_i c_i(z^c) = 0 \tag{48}$$

Futures contracts only induce wealth transfers; quasilinearity of spot goods preferences implies that the spot market outcomes we analyzed in Section 4 are unaffected by futures market outcomes. Since futures payments are z-measurable, Lemma 1 applies, and i's effective z-measurable wealth when purchasing c_i net contracts is simply her z-certainty-equivalent spot market wealth plus her futures payoffs:

$$\bar{W}_i + (z - z^c)\, c_i \tag{49}$$

with $\bar{W}_i$ defined in (38). Recall that we defined Δ_i as the linear exposure of i's z-CE wealth $\bar{W}_i$ to the spot price z. From (49), buying c_i contracts adjusts this linear exposure to:

$$\Delta_i + c_i = \mu_i + \psi(\beta_i - \alpha_i \nu_i) + c_i \tag{50}$$

Market clearing (48) then implies that these delta adjustments sum to zero across consumers, leaving society's aggregate delta unchanged:

$$\sum_i (\Delta_i + c_i) = \sum_i \Delta_i$$

Futures contracts are thus effectively a way for consumers to trade delta risk.

Building on Lemma 2, we first characterize i's risk-neutral distribution when purchasing c_i contracts, which we call $\mathbb{Q}_i^{c_i}$.

Lemma 3. *Suppose the following integrability condition holds:*

$$\Gamma_i > -\frac{1}{2\alpha_i \sigma_p^2} \tag{51}$$

Then, for any contract position c_i and price z^c, i's risk-neutral distribution $\mathbb{Q}_i^{c_i}$ is Gaussian, with mean and variance:

$$E_{\mathbb{Q}_i^{c_i}}(z) = -\alpha_i(\Delta_i + c_i)\, V_i \tag{52}$$

$$V_i \equiv Var_{\mathbb{Q}_i^{c_i}}(z) = \frac{1}{\frac{1}{\sigma_p^2} + 2\alpha_i \Gamma_i} \tag{53}$$

with Δ_i, Γ_i defined in (40) and (41) of Definition 2.

Proof. The wealth in (49) implies a risk-neutral distribution:

$$f_{\mathbb{Q}_i^{c_i}}(z) \propto U_i'\left(\bar{W}_i + (z - z^c)\, c_i\right) f_{\mathbb{P}}(z) \tag{54}$$

where we have used simply that, by the definition of the Radon-Nikodym derivative, the density of $\mathbb{Q}_i^{c_i}$ is proportional to the density of $\mathbb{P}$ multiplied by $\frac{d\mathbb{Q}_i^{c_i}}{d\mathbb{P}}$, which is the consumer's marginal utility. Now, for $f_{\mathbb{P}}(z)$, (27) implies that z is normally distributed with mean 0 and variance σ_p^2. From (38) and our assumption of CARA utility, marginal utility is:

$$U_i'\left(\bar{W}_i + (z - z^c)\, c_i\right) = \alpha_i \exp\left[-\alpha_i\left(A_i + \Delta_i z + \Gamma_i z^2 + (z - z^c)\, c_i\right)\right]$$

The RHS of (54) is thus:

$$f_{\mathbb{Q}_i^{c_i}}(z) \propto \exp\left(-\frac{z^2}{2\sigma_p^2} - \alpha_i\left(\Delta_i z + \Gamma_i z^2 + c_i\,(z - z^c)\right)\right) \tag{55}$$

Whenever (51) holds, (55) is a *Gaussian* density, with mean and variance (52) and (53) respectively. When (51) does not hold, the coefficient on z^2 in (55) is nonnegative. Thus, the $\mathbb{P}$-expectation of marginal utility is infinite; since under CARA utility $U_i'(W_i) = -\alpha_i U_i(W_i)$, this implies expected utility is $-\infty$. □

Lemma 3 represents the core analytical "trick" in our paper, so we pause to summarize the steps we have taken so far. First, under the quadratic goods-money preferences in (13), spot wealth W_i^{Spot} (20), or (30) in terms of z, is a quadratic function of spot prices z and an idiosyncratic η_i term. Second, from Lemma 1, CARA utility and joint normality of z, η_i imply that W_i^{Spot} collapses into the z-contingent "certainty equivalent" $\bar{W}_i$, which is a quadratic function of z, with no idiosyncratic term. Third, CARA utility implies that any z-quadratic wealths combine with the Gaussian physical probability measure $\mathbb{P}$ to create a Gaussian risk-neutral distribution $\mathbb{Q}_i^{c_i}$, whose mean and variance depend on the $\mathbb{P}$-variance σ_p^2; the polynomial "delta" and "gamma" coefficients of $\bar{W}_i$, Δ_i, Γ_i; i's risk aversion α_i; and any contracts c_i that i purchases.[5] This convenient property allows us to straightforwardly solve for futures contract demand.

[5]The "completing the square" approach in (55) is also isomorphic to Bayesian inference with normal priors, as well as the Kalman filter.

Proposition 3. *If (51) holds for i, i's contract demand at price z^c is:*

$$c_i\left(z^c\right) = -\mu_i + \psi\left(\nu_i\alpha_i - \beta_i\right) - \frac{z^c}{\alpha_i V_i} \tag{56}$$

with V_i defined in (53). If (51) holds for all i, the equilibrium contract price is:

$$z^c_{eq} = \psi\frac{\sum_i \nu_i\alpha_i + \sum_i \kappa_i}{\sum_i \frac{1}{\alpha_i V_i}} \tag{57}$$

Proof. The futures contract pays $f\left(z\right) = z$. Thus, (47) of Lemma 2 implies:

$$E_{\mathbb{Q}_i^{c_i}}\left[z\right] = z^c$$

Using (52) to substitute for $E_{\mathbb{Q}_i^{c_i}}\left[z\right]$, we require:

$$-\alpha_i\left(\Delta_i + c_i\right)V_i = z^c \tag{58}$$

Solving for c_i and substituting for Δ_i using (40), we obtain (56).

To solve for the equilibrium futures price z^c_{eq}, we sum futures demand across consumers and impose market clearing. From (58), it is instructive to write the market clearing condition as:

$$z^c_{eq}\sum_i \frac{1}{\alpha_i V_i} = -\sum_i \Delta_i \tag{59}$$

Intuitively, z^c_{eq} depends on $\sum_i \Delta_i$, which can be thought of as the economy's aggregate "delta", or linear exposure to the spot price z, divided by the harmonic sum of the product of risk aversions α_i and $\mathbb{Q}_i^{c_i}$-variances V_i, defined in (53). Substituting for Δ_i using (40), we have:

$$z^c_{eq}\sum_i \frac{1}{\alpha_i V_i} = -\sum_i\left(\mu_i + \psi\left(\beta_i - \alpha_i\nu_i\right)\right) \tag{60}$$

In (5), we normalized $\sum_i \mu_i = 0$, and from (33) we have $\sum_i \beta_i = -\sum_i \kappa_i$. Thus, we have:

$$z^c_{eq}\sum_i \frac{1}{\alpha_i V_i} = \psi\left(\sum_i \nu_i\alpha_i + \sum_i \kappa_i\right)$$

which gives (57). $\square$

Expression (56) shows that i's futures demand intercept contains three terms: $-\mu_i$, which reflects i's expected trade volume in spot markets; $-\psi\beta_i$, which measures correlations between i's inventory shock and the spot price; and $\psi\nu_i\alpha_i$, which is the product of i's risk aversion α_i

and the variance of the idiosyncratic component of her inventory shock, ν_i. We proceed to discuss how each term affects futures trade and prices.

6.1 Expected Spot Market Trades

Suppose $\psi = 0$. From (40), we have:

$$\Delta_i = \mu_i$$

that is, i's natural price exposure is just her expected inventory shock μ_i, which from (18) is the negative of her expected spot market purchase quantity. From (57), there is no equilibrium futures risk premium, $z^c_{eq} = 0$ and $p^c = \psi = 0$; from (56), we have:

$$c_i\left(z^c_{eq}\right) = -\mu_i$$

Futures trade is thus driven by the classical force of hedging expected trade volumes. In the buyer-seller example of Subsection 4.3, the buyer goes long a unit of the futures contract and the seller goes short a unit. This is welfare-improving because it eliminates both sides' directional exposure to spot prices, which shows up as the μ_i term in (40).

Keynes (1930) argues that futures risk premia may be driven by "hedging pressure", for example from farmers who hedge expected crop sales using short futures positions, leading to "normal backwardation". Hirshleifer (1990) challenges this intuition in an effectively complete general-equilibrium model, showing that producers' hedging positions can change sign without changing the futures risk premium. As an intuition, Hirshleifer writes: "consumers and producers are *mutually* hedging by taking the opposite sides of the futures transaction."

Our incomplete-market model preserves this "mutual hedging" logic for expected-trade-driven price exposures, and gives it a clean technical basis. Consumers' natural deltas move one-for-one with expected trade volumes (μ_i): expected sellers (buyers) have positive (negative) price exposures. These μ_i differences generate futures *trade*, but since $\sum_i \mu_i = 0$,[6] they contribute nothing to society's aggregate delta. Intuitively, spot markets always clear, so society as a whole is neither a buyer nor a seller of the spot good in expectation. Thus, beyond the consumer-producer example analyzed by Hirshleifer, expected trade volumes can always be mutually hedged in futures markets without creating aggregate pressure on the equilibrium futures price.[7]

[6] Once again, while we assumed $\sum_i \mu_i = 0$ in (5), this is purely a normalization. Appendix A.1 shows that any model in which aggregate inventory shocks are not mean-zero can be written as a model with mean-zero inventory shocks, where we simply redefine ψ to absorb the average inventory shock.

[7] As Hirshleifer observes, hedging pressure can create futures risk premia when there are participation constraints; exogenous participation constraints are used to create futures risk premia in many models

6.2 Inventory Shock Hedging

When $\psi \neq 0$, consumers' contract demand contains a $-\beta_i\psi$ term. Recall that β_i is the projection coefficient in (31) of Definition 1, capturing the correlation between i's inventory shock ϵ_i and the spot price, determined by (32):

$$\beta_i = -\left(\sum_j \kappa_j\right)\frac{\sigma_i^2}{\sum_j \sigma_j^2}$$

β_i is thus weakly negative, and $-\beta_i\psi$ always has the same sign as ψ. Intuitively, this term captures the fact that the spot price is informative about i's inventory shock, so i can partially hedge inventory risk using futures contracts.

Consider the maker-taker example of Subsection 4.2. Expected shocks μ_T, μ_M are both 0, and basis risk variances ν_T, ν_M are zero because the spot price is perfectly correlated with ϵ_T. The only driver of contract trade is the $\psi\beta_i$ term; we have:

$$\beta_T = -\kappa_M,\ \beta_M = 0$$

Plugging (57) into (56), each agent's equilibrium linear price exposure, $c_i + \Delta_i$, satisfies:

$$c_i\left(z^c_{eq}\right) + \Delta_i = c_i\left(z^c_{eq}\right) + \beta_i\psi = -\psi\kappa_M\left(\frac{\frac{1}{\alpha_i V_i}}{\sum_i \frac{1}{\alpha_i V_i}}\right) \tag{61}$$

Hence, in equilibrium, the liquidity taker T is long futures and the market maker M is short.

We noted in Subsection 4.2 that spot market outcomes leave T holding her entire linear "inventory risk" $x_T\psi$, whereas the first-best outcome involves T sharing this term with M. Somewhat surprisingly, this example shows that futures markets allow T and M to share this inventory risk: T's long futures position pays out when prices are high and x_T is low, reducing her exposure to her own inventory, and correspondingly transferring some exposure to M. Exposure transfers are particularly straightforward in this example because prices are a perfect signal for T's inventory shock.

The $\psi\beta_i$ components of consumers' deltas sum to a simple aggregate delta position of $-\psi\sum_i \kappa_i$. Technically, $\sum_i \beta_i$ reflects the coefficient from regressing the aggregate inventory shock $\sum_i \epsilon_i$ on the spot price; from (27), this is equal to the slope of aggregate demand, $-\sum_i \kappa_i$. Intuitively, society's aggregate delta, $\sum_i \Delta_i$, represents the marginal relationship of social aggregate wealth to the spot price; a one-unit increase in the spot price implies a

(Hirshleifer, 1988*a*,*b*; Gorton, Hayashi and Rouwenhorst, 2013; Acharya, Lochstoer and Ramadorai, 2013; Goldstein, Li and Yang, 2014; Goldstein and Yang, 2022).

decrease of $\sum_i \kappa_i$ in aggregate inventory, which is marginally worth $\psi \sum_i \kappa_i$ dollars.

Expression (57) shows that this aggregate delta creates a futures risk premium. If we ignore the $\nu_i \alpha_i$ term – as in the M, T example where ν_T, ν_M are both zero – we can write (57) as:

$$\frac{p^c_{eq} - \psi}{\psi} = \frac{\sum_i \kappa_i}{\sum_i \frac{1}{\alpha_i V_i}} \tag{62}$$

That is, the futures price p^c contains a proportional markup over the $\mathbb{P}$-fair price ψ, which depends on the aggregate delta $\sum_i \kappa_i$ divided by the harmonic sum of risk-aversion-weighted $\mathbb{Q}$-variance $\alpha_i V_i$. Intuitively, since long futures positions hedge negative aggregate inventory shocks, society in the aggregate would like to be long futures at the fair price ψ; the markup (62) must be positive to lower futures demand sufficiently to clear markets. The denominator of (62), $\sum_i \frac{1}{\alpha_i V_i}$, is the futures *demand slope*: it is driven by risk aversion and $\mathbb{Q}$-variances, as in the familiar problem of optimal portfolio choice in normal-exponential mean-variance models.

Returning to the M, T example, substituting (16) into first-best wealth (22), we can write:

$$G^*_T = C_T + \frac{\alpha_T^{-1}}{\alpha_T^{-1} + \alpha_M^{-1}} \left(-\psi \kappa_M (p - \psi) - \frac{1}{2} \kappa_M (p - \psi)^2 \right) \tag{63}$$

The linear price exposure in first-best wealth (63) looks quite similar to contract-induced wealth (61); the difference is that there is a $\frac{1}{\alpha_i V_i}$ weighting term in futures equilibrium, versus a $\frac{1}{\alpha_i}$ term in the first-best. In the following section, we will show that variance swaps "fix" this issue, aligning equilibrium weights with the first-best outcome.

6.2.1 Inventory Risk and Optimal Hedging

Suppose a group of farmers, F, receive a random harvest, which they inelastically sell to a market with some aggregate demand slope. The crop is valuable on average, so $\psi > 0$. Farmers have $\mu_F > 0$, since they are expected sellers; this tends to push them towards shorting futures. We assume all inventory shocks accrue to farmers, so their variance share is $\frac{\sigma_F^2}{\sum_i \sigma_i^2} = 1$, and basis risk is $\nu_F = 0$. Counterintuitively, since farmers tend to have large negative β_F, inventory risk pushes farmers towards being *long* futures.

To see how these forces interact, suppose contracts are $\mathbb{P}$-fairly priced, so $z^c = 0$, that is, $p^c = \psi$. From (56), contract demand is:

$$c_F(0) = -\mu_F - \psi \beta_F \tag{64}$$

Now, $\frac{\sigma_F^2}{\sum_i \sigma_i^2} = 1$ implies $\beta_F = -\sum_{i=1}^N \kappa_i$. Thus, c_F is positive if:

$$\frac{\psi}{\mu_F} \sum_{i=1}^{N} \kappa_i > 1 \tag{65}$$

Recall $\sum_{i=1}^N \kappa_i$ is the spot demand slope, $\psi = p^c$ is the contract price, and μ_F is the expected quantity farmers sell; farmers are thus long or short futures depending on whether the *demand elasticity* for crops is greater than 1.

McKinnon (1967) and Rolfo (1980) analyze farmers' optimal hedging problem, deriving conditions analogous to (65). But these papers do not close the futures market: who takes the other side of the farmers' hedge, what futures premium clears the market, and what are the welfare consequences?

Our model shows that (65) reflects the tension between the expected-trade hedging (μ_i) and inventory-risk sharing ($\psi\beta_i$) roles of futures contracts. The latter force contributes risk premia, whereas the former always nets out across agents. Either motive can dominate futures trade.

6.3 Basis Risk

The final term in contract demand (56) involves *basis risk*, $\psi\nu_i\alpha_i$. Since ν_i and α_i are nonnegative, this term always shares the sign of ψ. Technically, spot wealth (34) has an idiosyncratic $(\psi + z)\,\eta_i$ term; Lemma 1 translates this into a z-CE wealth penalty proportional to $(\psi + z)^2\,\nu_i\alpha_i$. Intuitively, inventory shocks create wealth risk which is greater when the spot price $(\psi + z)$ is large in absolute value; when ψ is positive, high spot prices increase wealth risk and thus decrease consumers' z-CE wealth.

As in the previous examples, heterogeneity in consumers' natural Δ_i exposures leads to gains from futures trade. Consider a simple example where $\psi > 0$, $\nu_i = \nu$ and $\beta_i = \beta$ are constants, and all agents have $\mu_i = 0$, but agents differ in their risk aversions α_i. Contract demand is then simply:

$$c_i\,(z^c) = \psi\,(\nu\alpha_i - \beta) - \frac{z^c}{\alpha_i V_i}$$

Intuitively, all agents suffer greater wealth losses from basis risk when prices are higher, generating implicit short delta exposure to spot prices. Agents with larger risk aversion α_i suffer more, in z-CE terms. In equilibrium, higher α_i agents thus take long futures positions, selling some of their wealth risk to lower-α_i agents, who are short futures in equilibrium.

The $\nu_i\alpha_i$ term is always positive, so it does not net out across agents. We can in general write the contract pricing condition, (57), as:

$$\frac{p^c_{eq} - \psi}{\psi} = \frac{\sum_i \nu_i \alpha_i + \sum_i \kappa_i}{\sum_i \frac{1}{\alpha_i V_i}} \tag{66}$$

Intuitively, futures risk premia can in general be thought of as resulting from the aggregate *wealth betas* of spot prices. This has two components: correlations between spot prices and aggregate inventory, captured by the $\sum_i \kappa_i$ term, and the fact that higher prices amplify basis risk and lower agents' z-CEs, captured by the $\sum_i \nu_i \alpha_i$ term. Both terms are always positive, implying that the futures risk premium always has the same sign as ψ.

If markets were complete and agents were able to eliminate the basis risk term $\sum_i \nu_i \alpha_i$, the numerator of (66) would be smaller; in our model, market incompleteness magnifies society's aggregate delta position, because high prices amplify the negative effects of inventory shocks. Market incompleteness also increases agents' $\mathbb{Q}$-variances V_i, by lowering Γ_i in (53), which reduces the denominator of (66). Therefore, market incompleteness always increases the magnitude of the futures risk premium.

To summarize, futures contracts allow agents to trade their natural directional exposure to spot prices, which we call Δ_i. Optimal risk sharing always involves sharing Δ_i according to agents' risk aversions, and their $\mathbb{Q}$-variances V_i in (53). Agents' Δ_is have three terms – the expected spot trade term, μ_i; the inventory-price correlation term, $\psi \beta_i$; and the spot-price-basis-risk correlation term, $\psi \nu_i \alpha_i$; in general, futures trade can be driven by all three of these forces.

7 Variance Swaps

Next, we add *variance swaps.* A consumer who purchases a unit of a variance swap pays some fixed price q^c in all future states, and receives $z^2 = (p - \psi)^2$.[8] Applying Lemma 1 and (38), a consumer i who purchases c_i futures contracts and d_i variance swaps at respective prices z^c, q^c attains effective z-measurable wealth:

$$\bar{W}_i + (z - z^c)\, c_i + \left(z^2 - q^c\right) d_i$$

[8] Clearly, any noncentered variance swap is payoff-equivalent to some combination of a centered variance swap and a futures contract position.

with $\bar{W}_i$ defined in (38). Analogously to the futures case, quadratic derivatives function by adjusting agents' quadratic exposure to prices, to:

$$\Gamma_i + d_i = \frac{\kappa_i}{2} + \beta_i - \frac{\alpha_i \nu_i}{2} + d_i$$

Variance swaps thus allow consumers to trade gamma risk – second-order price risk – leaving society's aggregate gamma unchanged. We first solve for i's $\mathbb{Q}_i^{c_i,d_i}$-measure.

Lemma 4. *For any futures and variance swap positions c_i, d_i and prices z^c, q^c, suppose d_i satisfies:*

$$\Gamma_i + d_i > -\frac{1}{2\alpha_i \sigma_p^2} \tag{67}$$

Then i's risk-neutral distribution $\mathbb{Q}_i^{c_i,d_i}$ is Gaussian, with mean and variance:

$$m_i \equiv E_{\mathbb{Q}_i^{c_i,d_i}}(z) = -\alpha_i (\Delta_i + c_i) V_i \tag{68}$$

$$V_i \equiv Var_{\mathbb{Q}_i^{c_i,d_i}}(z) = \frac{1}{\frac{1}{\sigma_p^2} + 2\alpha_i (\Gamma_i + d_i)} \tag{69}$$

with Δ_i, Γ_i defined in (40) and (41) of Definition 2.

Proof. As in the proof of Lemma 3, combining CARA utility with the normal distribution of the spot price z yields a density for $\mathbb{Q}_i^{c_i,d_i}$ that satisfies:

$$f_{\mathbb{Q}_i^{c_i,d_i}}(z) \propto \exp\left(-\frac{z^2}{2\sigma_p^2} - \alpha_i \left(\Delta_i z + \Gamma_i z^2 + c_i (z - z^c) + d_i \left(z^2 - q^c\right)\right)\right) \tag{70}$$

This again describes a Gaussian density with mean and variance (68) and (69). Intuitively, the $\mathbb{Q}_i^{c_i,d_i}$-density is normal as long as the exponent of marginal utility is at most a second-order polynomial in z; variance swaps simply adjust the coefficient on z^2, and thus the normality of $\mathbb{Q}_i^{c_i,d_i}$ is preserved. Analogous to (51) of Lemma 3, (67) guarantees that expected marginal utility is finite. □

We then use risk-neutral pricing to solve simultaneously for futures and variance swap demand and equilibrium prices.

Proposition 4. *Let:*

$$V(q^c, z^c) \equiv q^c - (z^c)^2 \tag{71}$$

be the "price of spot price variance" implied by q^c, z^c. If $V(q^c, z^c) > 0$, i's contract demand *exists, and is:*

$$d_i\left(q^c, z^c\right) = \frac{1}{2\alpha_i}\left(\frac{1}{V} - \frac{1}{\sigma_p^2}\right) - \Gamma_i \tag{72}$$

$$c_i\left(q^c, z^c\right) = -\mu_i + \psi\left(\nu_i\alpha_i - \beta_i\right) - \frac{z^c}{\alpha_i V} \tag{73}$$

with Δ_i, Γ_i defined in (40) and (41) of Definition 2. An equilibrium in the market for futures and variance swaps exists whenever:

$$\sum_i \frac{1}{\alpha_i} > \sigma_p^2 \sum_i \left(\kappa_i + \alpha_i\nu_i\right) \tag{74}$$

Equilibrium prices q_{eq}^c, z_{eq}^c satisfy the system of equations:

$$\left(\frac{1}{V\left(q_{eq}^c, z_{eq}^c\right)} - \frac{1}{\sigma_p^2}\right)\sum_i \frac{1}{\alpha_i} = -\sum_i \left(\kappa_i + \alpha_i\nu_i\right) \tag{75}$$

$$z_{eq}^c = \psi \frac{\sum_i \nu_i\alpha_i + \sum_i \kappa_i}{\frac{1}{V\left(q_{eq}^c, z_{eq}^c\right)}\sum_i \frac{1}{\alpha_i}} \tag{76}$$

Proof. See Appendix A.6. □

7.1 Market Makers and Liquidity Takers

In the real world, variance is traded in *option markets*. Option pricing is one of the foundational results in asset pricing (Black and Scholes, 1973) and the subject of a large and ongoing literature. But what role, if any, do options on goods prices play in improving risk-sharing and social welfare? Futures contracts are intuitively valuable because they allow expected buyers and sellers to mutually hedge directional price risk; it is less obvious why agents would ever have offsetting price *variance* exposures to trade in option markets.[9]

We give a simple answer based on the example of Subsection 4.2. We showed that the "liquidity taker" has natural short exposure to inventory variance, and the "market maker" is naturally long variance. Through (32) and (41), this shows up in their gamma exposures to price variance:

$$\Gamma_T = \beta_T = -\kappa_M,\ \Gamma_M = \frac{\kappa_M}{2} \tag{77}$$

[9] As we discussed in Subsection 1.1, related work by Brennan and Cao (1996) and Chabakauri, Yuan and Zachariadis (2022) provides *information-based* rationales for volatility derivative trade. Their models analyze financial asset markets without explicit spot markets; volatility trade arises from heterogeneous valuations of quadratic payoffs induced by asymmetric information. We abstract from informational frictions, and instead study how variance swaps redistribute second-order price exposures generated by spot-market trade.

where we have used that, since z is perfectly correlated with T's inventory, $\nu_T = \nu_M = 0$. Exactly analogous to how buyers and sellers neutralize their offsetting *directional* exposures to price risk through futures contracts, M and T can neutralize their offsetting exposures to price *variance* risk using variance swaps.

Intuitively, suppose T is long $\frac{\kappa_M}{2}$ variance swaps, and M is correspondingly short. M then perfectly eliminates her price variance exposure: her short volatility position pays off in states of the world where the aggregate inventory shock is near 0, so her liquidity provision profits are low. T eliminates half of her negative exposure, since her long swap position pays off when inventory shocks are large, partially offsetting her spot market wealth losses. This trade is thus clearly Pareto-improving.

M and T cannot fully offset their price variance exposures because, from (77), society's aggregate gamma is negative and equal to half the market maker's demand slope, $-\frac{\kappa_M}{2}$. Analogous to the futures case, variance swap markets redistribute society's aggregate gamma risk according to agents' risk capacities. Combining (72) and (75), in market equilibrium, each agent's final gamma position $d_i + \Gamma_i$ satisfies:

$$d_T\left(q^c_{eq}, z^c_{eq}\right) + \Gamma_T = -\frac{\alpha_T^{-1}}{\alpha_T^{-1} + \alpha_M^{-1}}\frac{\kappa_M}{2}, \; d_M\left(q^c_{eq}, z^c_{eq}\right) + \Gamma_M = -\frac{\alpha_M^{-1}}{\alpha_T^{-1} + \alpha_M^{-1}}\frac{\kappa_M}{2} \tag{78}$$

After substituting inventory shocks for prices in (78) using (16), the resulting expressions are exactly the efficient quadratic exposures described in (22) and (23). Variance swap markets thus allow T and M to efficiently share price variance risk.

In this example, variance swaps accomplish the division of labor between physical production and financial risk-bearing discussed in Bernstein (1996) and Shiller (2012). It is socially efficient for T to sell her entire inventory to M in spot markets; this causes M to be naturally long volatility, and T to be naturally short. But variance swap markets fix this, by first neutralizing any offsetting volatility exposures, and then allocating the residual price volatility risk to agents according to their inverse risk aversions.[10]

[10]This mechanism is distinct from the dealer-inventory channel in Gârleanu, Pedersen and Poteshman (2009). In their index-option application, end-users are empirically net long options, especially out-of-money puts, and dealers are correspondingly net short. However, end-user demand, and hence dealer inventory, is a primitive input to their pricing model. In our setting, derivative positions are equilibrium outcomes: T is naturally short volatility and M naturally long volatility because of their roles in the spot market, leading T to buy and M to sell variance swaps in equilibrium.

7.2 Markets for Gamma

In the general case, (72) can be interpreted as saying that each consumer sells her natural Γ_i exposure, then purchases an amount of contracts depending on the gap between the $\mathbb{P}$-fair price of a variance swap, σ_p^2, and the market price V of the swap, as well as i's risk aversion α_i. (94) shows that the equilibrium *variance risk premium*, $\frac{1}{V} - \frac{1}{\sigma_p^2}$, depends on society's aggregate Γ_i and aggregate risk aversion. This simplifies substantially, due to the identity (33): the β_i terms in (41) cancel, leaving us with the RHS of (75), $-\sum_i (\kappa_i + \alpha_i \nu_i)$.

Intuitively, in our model, society has negative gamma in the aggregate – concave exposure to the spot price – for two reasons. First, demand for the spot good slopes downwards – society in the aggregate has declining marginal utility for the good. Social dollar-valued wealth is thus concave in the aggregate inventory shock $\sum_i \epsilon_i$ – negative inventory shocks hurt more than positive shocks help – and thus is also concave in prices, which are linear in inventory.[11]

Second, in (36) of Lemma 1, basis risk enters consumers' z-CE wealth multiplicatively with the squared price $(\psi + z)^2$; thus, basis risk is more costly when the spot price is larger in absolute value. This term increases the magnitude of consumers' gamma exposures through (41), and thus increases society's aggregate gamma exposure and thus the equilibrium variance risk premium.

7.3 Basis Risk-Driven Variance Trading

Since basis risk $\alpha_i \nu_i$ also appears in agents' spot gamma exposures, basis risk can also drive variance swap trading. Suppose we have N agents, with $\mu_i = 0$, $\psi = 0$, and identical κ and

[11] Note that, from (11), aggregate wealth is:

$$W^*(\boldsymbol{x}) = \psi \sum_{i=1}^{N} x_i - \frac{\left(\sum_{i=1}^{N} x_i\right)^2}{2\sum_{i=1}^{N} \kappa_i}$$

Substituting for $\sum_{i=1}^{N} x_i$ using (16), we have:

$$W^*(\boldsymbol{x}) = -\psi \left(\sum_{i=1}^{N} \kappa_i\right)(p - \psi) - \frac{\left(\sum_{i=1}^{N} \kappa_i\right)(p-\psi)^2}{2}$$

Hence, society's aggregate wealth exposure to $(p-\psi)^2$ is in fact exactly $-\frac{\sum_{i=1}^{N} \kappa_i}{2}$.

inventory shock variance σ^2. From (32), we have:

$$\beta_i = \beta = -\kappa,\ \nu_i = \nu = \sigma^2 \left(\frac{N-1}{N}\right)$$

Thus, the Greeks are:

$$\Delta_i = 0, \Gamma_i = -\frac{\kappa}{2} - \frac{\alpha_i}{2}\sigma^2 \left(\frac{N-1}{N}\right)$$

Thus, high α_i agents have comparatively large negative gammas, through the $\alpha_i \nu_i$ term in (41), whereas optimal risk-sharing requires these consumers to hold comparatively small shares of total gamma, as in (78) of the example above. Society's total gamma exposure is:

$$\sum_{i=1}^{N} \Gamma_i = -\frac{N\kappa}{2} - \frac{\sum_{i=1}^{N} \alpha_i}{2}\sigma^2 \left(\frac{N-1}{N}\right)$$

Applying (72) and (75), agents' equilibrium gamma exposures are:

$$d_i \left(q_{eq}^c, z_{eq}^c\right) + \Gamma_i = \frac{\frac{1}{\alpha_i}}{\sum_i \frac{1}{\alpha_i}} \left(-\frac{N\kappa}{2} - \frac{\sum_{i=1}^{N} \alpha_i}{2}\sigma^2 \left(\frac{N-1}{N}\right)\right)$$

Equilibrium thus involves high-alpha consumers "buying variance" from low-alpha consumers, again moving from consumers' natural spot market Greeks to their efficient Greeks. Note that, in this case, variance swaps do not eliminate basis risk; rather, they allow agents to trade the price-dependent impact of basis risk – the fact that the z-CE cost of basis risk in (36) is proportional to α_i and the price z^2 – thus loading this implicit squared price risk preferentially onto less risk-averse agents.

7.4 Futures Trade and Pricing with Variance Swaps

Variance swap trading influences futures demand by changing agents' $\mathbb{Q}$-variances. This changes the slope of futures demand – the z^c coefficient in (73). This does not qualitatively change futures market outcomes: the basic patterns in the examples we analyzed in the previous section persist when variance swaps are traded alongside futures. However, quantitatively, from (73) and (76) of Proposition 4, and using the definition of Δ_i in (40), agents' equilibrium delta exposures satisfy:

$$c_i \left(q_{eq}^c, z_{eq}^c\right) + \Delta_i = \left(\sum_j \Delta_j\right) \left(\frac{\frac{1}{\alpha_i}}{\sum_j \frac{1}{\alpha_j}}\right)$$

Consider again the maker-taker example of Subsection 4.2; analogous to our derivations in Subsection 6.2 in the futures-only case, we find:

$$c_i\left(q_{eq}^c, z_{eq}^c\right) + \Delta_i = c_i\left(q_{eq}^c, z_{eq}^c\right) + \beta_i\psi = -\psi\kappa_M\left(\frac{\frac{1}{\alpha_i}}{\sum_i \frac{1}{\alpha_i}}\right) \tag{79}$$

Expression (79) matches the linear exposure in first-best wealth, as we wrote in (63). Intuitively, equilibrium outcomes with only linear futures contracts differ from first-best risk-sharing because agents' $\mathbb{Q}$-variances do not agree, adding an extra $\mathbb{Q}$-variance weight, $\frac{1}{\alpha_i V_i}$, in delta-sharing through futures contracts in equilibrium. But variance swap trading perfectly aligns agents' $\mathbb{Q}$-variances: this is exactly the equilibrium condition (75) in Proposition 4. Thus, variance swap trading causes the common V_i term to drop out of (79), leading linear contracts to also implement exactly the inverse-risk-aversion formula that characterizes first-best risk-sharing. We will generalize this result in Proposition 5 below, showing that futures and variance swaps achieve price-contingent completeness.

An interesting technical result is that, while variance swaps influence futures trade, they happen not to change equilibrium futures pricing.

Lemma 5. *The futures risk premium is identical with or without variance swaps, and is equal to:*

$$z_{eq}^c = \psi\frac{\sum_i \kappa_i + \sum_i \nu_i\alpha_i}{\frac{1}{\sigma_p^2}\sum_i \frac{1}{\alpha_i} - \left(\sum_i \kappa_i + \sum_i \alpha_i\nu_i\right)}$$

Proof. See Appendix A.7. □

Intuitively, the futures risk premium is pinned down by society's aggregate delta position divided by the aggregate slope of futures demand. In our CARA-normal-quadratic setting, each agent's demand slope is affine in her post-trade gamma exposures. Variance swaps redistribute gamma across agents without changing its net supply, so they do not influence the aggregate demand slope, and thus the futures risk premium. We view this result as a convenient mathematical consequence of our functional form assumptions, rather than a robust economic prediction.

7.5 Price-Contingent Completeness

Classical *complete markets* implement allocations that cannot be Pareto-improved through state-contingent transfers. Price-derivative markets are incomplete because basis risks remain unspanned. Nevertheless, we show that equilibrium with futures and variance swaps implements an allocation that cannot be Pareto-improved through *price-contingent* wealth transfers. We call this property *price-contingent completeness.*

Definition 3. A state-contingent wealth allocation satisfies *price-contingent completeness* if it is impossible to Pareto-improve upon the allocation using budget-balanced *price-contingent* wealth transfers. Formally, suppose the planner implements arbitrary z-contingent wealth transfers:

$$t_1(z), t_2(z), \ldots, t_N(z)$$

satisfying budget balance for all z:

$$\sum_i t_i(z) = 0$$

An allocation satisfies price-contingent completeness if no such choice of $t_i(z)$ leads to a Pareto improvement in agents' expected utilities.

Price-contingent completeness is only meaningful in incomplete-markets settings: if the spot price spans the state space, price-contingent completeness is trivially equivalent to first-best risk sharing.

Proposition 5. *The equilibrium allocation with futures and variance swaps satisfies price-contingent completeness.*

Proof. By Lemma 1, the z-CE $\bar{W}_i$ continues to represent agent i's preferences after any z-measurable transfer. Define i's equilibrium z-contingent wealth, including her futures and variance swap positions, as:

$$\bar{W}_i^{FinEq}(z) \equiv \bar{W}_i + \left(z - z_{eq}^c\right) c_i\left(q_{eq}^c, z_{eq}^c\right) + \left(z^2 - q_{eq}^c\right) d_i\left(q_{eq}^c, z_{eq}^c\right) \tag{80}$$

The planner's problem therefore reduces to a standard complete-markets risk-sharing problem on the state space generated by z. Thus, an allocation satisfies price-contingent completeness if and only if the Borch (1962) risk-sharing condition holds for z-CE wealths. Thus, we seek positive Pareto weights λ_i such that, for every pair of agents i, j,

$$\lambda_i U_i'\left(\bar{W}_i^{FinEq}(z)\right) = \lambda_j U_j'\left(\bar{W}_j^{FinEq}(z)\right)$$

for $\mathbb{P}$-almost every z. In asset pricing terms, write $\mathbb{Q}_i$ for i's equilibrium risk-neutral measure; normalizing each agent's marginal utility to integrate to one, the Borch condition becomes:

$$\frac{d\mathbb{Q}_i}{d\mathbb{P}} = \frac{d\mathbb{Q}_j}{d\mathbb{P}}$$

for $\mathbb{P}$-almost every z. Thus, it suffices to show that all agents have the same risk-neutral measure in equilibrium.

The equilibrium first-order conditions for futures and variance swaps respectively imply

$$E_{\mathbb{Q}_i}[z] = z^c_{eq}, \qquad E_{\mathbb{Q}_i}[z^2] = q^c_{eq}$$

for all i. Hence, all agents' risk-neutral measures have the same first and second moments. From Lemma 4, each $\mathbb{Q}_i$ is Gaussian. Because a Gaussian distribution is fully determined by its first and second moments, all agents' risk-neutral measures coincide. The Borch condition thus holds at the equilibrium allocation, so no budget-balanced z-contingent transfers can generate a Pareto improvement. □

Intuitively, the result follows from the two-dimensional structure of agents' equilibrium pricing measures. Each $\mathbb{Q}_i$ is Gaussian, and is therefore identified by its first two moments, $E_{\mathbb{Q}_i}[z]$ and $E_{\mathbb{Q}_i}[z^2]$. Futures and variance swaps trade precisely these two statistics, so equilibrium aligns agents' entire $\mathbb{Q}$-measures and hence their normalized marginal utilities. Traded contracts thus need not span every z-contingent payoff: they only need to trade the two statistics that pin down the associated Gaussian distribution.[12]

Price-contingent completeness is related to, but distinct from, the notion of *conditional Pareto efficiency* in Brennan and Cao (1996) and the analogous notion of *effective completeness* in Chabakauri, Yuan and Zachariadis (2022). In these models, private information gives informed and uninformed investors different posteriors over terminal payoff states. Trading the underlying aligns investors' risk-neutral means but leaves an information-driven wedge in their risk-neutral variances; a quadratic claim allows investors to trade this variance exposure until their information-contingent risk-neutral measures coincide.

Our result shares the mathematical structure that linear and quadratic payoffs can implement the relevant Pareto frontier, similarly exploiting CARA preferences and an exponential-quadratic structure. However, the underlying risk-sharing problem in our setting is different. Agents trade before inventory shocks are realized and share a common prior over $\boldsymbol{x}$, but anticipate that their spot-market wealths $W_i^{Spot}(\boldsymbol{x})$ depend on those shocks. Price derivatives allow traders to share the price-measurable component of these exposures; price-contingent completeness occurs when the normalized marginal utilities of agents' z-CE wealths are aligned, implying that the planner cannot further Pareto-improve outcomes

[12] Since futures and variance swaps align agents' $\mathbb{Q}$-measures in equilibrium, we also attain a standard set of complete-market asset-pricing results: if we add any additional z-derivative satisfying certain boundedness conditions, there exists an equilibrium with no trade in the z-derivative, where the equilibrium price is pinned down by consumers' common $\mathbb{Q}$-measure.

through price-contingent transfers.[13]

Price-contingent completeness differs from the *constrained efficiency* notion in Geanakoplos and Polemarchakis (1986), which restricts the planner to reallocating existing financial assets. Our planner, in contrast, can choose arbitrary transfer functions that are contingent on the spot price z.

7.6 Welfare

Using Proposition 5, we will derive analytical expressions for welfare with futures and variance swaps. We will measure welfare using *certainty equivalents*.

Definition 4. i's certainty equivalent for a random wealth W_i is:

$$CE_i \equiv U_i^{-1}\left[E\left[U_i\left(W_i\right)\right]\right] = -\frac{1}{\alpha_i}\log\left[E\left[\exp\left(-\alpha_i W_i\right)\right]\right] \tag{81}$$

We utilize a useful aggregation property of CEs under CARA utility.

Lemma 6. *Suppose an arbitrary random total wealth W is optimally shared among agents, that is, (12) of Proposition 1 holds. Then the sum of agents' certainty equivalents is:*

$$\sum_{i=1}^{N} CE_i = -\frac{1}{\alpha^{agg}}\log E\left[\exp\left(-\alpha^{agg} W\right)\right] \tag{82}$$

where

$$\alpha^{agg} \equiv \left(\sum_{i=1}^{N}\alpha_i^{-1}\right)^{-1} \tag{83}$$

Proof. This follows immediately from Wilson (1968); we give a formal derivation in Appendix A.8. □

Formally, Lemma 6 states that, under optimal risk sharing, the N agents behave like a single representative agent with risk aversion (83). The representative agent's CE, (82), is a useful measure of aggregate social welfare delivered by random wealth W under perfect

[13] There is also a timing difference between our setting and those of Brennan and Cao (1996) and Chabakauri, Yuan and Zachariadis (2022). Their efficiency and risk-sharing criteria are conditional on the information available to investors after private information and asset prices are observed; information arriving before risk sharing is complete can destroy ex ante risk-sharing opportunities (Hirshleifer, 1971). In our model, agents are symmetrically informed and trade before inventory shocks are realized, so price-contingent completeness is an ex ante notion based on their common prior over future states, analogous to Borch (1962) and Wilson (1968).

risk-sharing: all agents could be made equally well off if W were replaced by a nonrandom constant equal to the RHS of (82).

Let CE^{FB} denote society's aggregate certainty equivalent under the first-best allocation $G_i^*(\boldsymbol{x})$ in Proposition 1, and let CE^{Equil} denote society's aggregate certainty equivalent under the derivatives-market equilibrium allocation $\bar{W}_i^{FinEq}(z)$ in the proof of Proposition 5.

Proposition 6. *Suppose:*

$$\sum_i \frac{1}{\alpha_i} > \sigma_p^2 \sum_i \kappa_i \tag{84}$$

Then agents' expected utility in the first-best outcome is finite; the aggregate CE under first-best risk sharing exists, and is:

$$CE^{FB} = \frac{\sum_{i=1}^N \kappa_i}{2}\psi^2 + \frac{1}{2\alpha^{agg}} \log\left(1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^N \kappa_i\right) - \frac{\psi^2 \sum_{i=1}^N \kappa_i}{2\left(1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^N \kappa_i\right)}. \tag{85}$$

Suppose the integrability condition (74) holds; then the aggregate CE with futures and variance swaps is:

$$CE^{Equil} =$$
$$\frac{\sum_{i=1}^N \kappa_i}{2}\psi^2 + \frac{1}{2\alpha^{agg}} \log\left(1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^N (\kappa_i + \alpha_i\nu_i)\right) - \frac{\psi^2 \sum_{i=1}^N (\kappa_i + \alpha_i\nu_i)}{2\left(1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^N (\kappa_i + \alpha_i\nu_i)\right)} \tag{86}$$

with α^{agg} defined in (83) of Lemma 6. Suppose CE^{FB} and CE^{Equil} both exist; if ψ^2 and σ_p^2 are $O(\epsilon)$, we have:

$$CE^{FB} - CE^{Equil} = \frac{\sum_{i=1}^N \alpha_i\nu_i}{2} E\left[p^2\right] + O\left(\epsilon^2\right) \tag{87}$$

Proof. See Appendix A.9. □

Expressions (85) and (86) are analytical expressions for first-best and equilibrium welfare. Intuitively, we can apply the aggregation result of Lemma 6 to the equilibrium case because we showed in Proposition 5 that agents engage in first-best risk sharing over z-CE wealths. Thus, (85) and (86) differ essentially in the addition of aggregated basis risk penalty terms $\sum_{i=1}^N \alpha_i\nu_i$.

Expression (87) approximates the *welfare costs of market incompleteness* – the gap between first-best and equilibrium CEs – when ψ^2 and σ_p^2 are small. Intuitively, in our construction of z-CEs, (36) shows that each agent i suffers a basis risk penalty of

$$\frac{\alpha_i \nu_i}{2} p^2$$

dollars. Summing across agents, and approximating p^2 with its expectation $E[p^2]$, we obtain (87).

The cost of market incompleteness thus depends on the extent to which prices span the space of risks, and the extent to which unspanned basis risk ν_i accrues to agents with high risk aversions α_i. These costs are zero in our examples in Subsections 4.2 and 4.3, since there is only a single inventory shock; costs are positive when multiple agents receive inventory shocks, as we analyzed in Subsections 6.3 and 7.3.

For completeness, in Appendix A.10, we derive an analytical expression for certainty-equivalent wealth in isolated spot markets. As we discussed above, price-linked derivative markets generate trade and thus welfare gains through a number of different economic channels; thus, the difference between welfare in isolated spot markets and welfare with price-linked derivatives does not reduce to an expression as economically intuitive as (87).

8 Discussion of Model Assumptions

Formally, we have analyzed a model of *general equilibrium with incomplete markets* (GEI). Here, we discuss the main restrictive assumptions we impose, and the roles they play in our analysis.

First, we assume in (3) that consumers have *quasilinear* preferences over goods and money.[14] This eliminates income effects, so wealth transfers across agents do not affect spot market demands; under our concave production technology, spot market equilibria are unique and unaffected by wealth transfers across agents. Our model thus inherits attractive properties of the *partial equilibrium* setting of Marshall (1920): allocative efficiency is naturally defined as the money-metric value of equilibrium outcomes, allowing us to decompose any deviation from optimal outcomes into allocative and risk-sharing components. Like classic partial-equilibrium models, our setting is less appropriate for studying "large" goods where income effects on demand are important.[15]

[14] Quasilinear or transferable-utility GEI models are relatively uncommon; examples include Dierker and Dierker (2010) and Dierker and Dierker (2012), who study incomplete-market production economies with a single firm rather than derivative trading or CARA-normal-quadratic risk sharing.

[15] A classic modelling question is when a good is "small" enough that income effects can be ignored; see, for example, Vives (1987). In the commodity-market applications we have in mind, this approximation is most natural for producers, intermediaries, and large industrial consumers, whose spot demands are

Second, we assume CARA utility; this drives the existence of transfer-independent z-CEs. From Lemma 1, a z-CE collapses basis risk into a certainty-equivalent monetary penalty, independent of arbitrary z-contingent transfers. As we discussed around Footnote 3, this is possible because of the level-invariance property of CARA utility. Under general concave utility functions, transfer-independent certainty equivalents would not be possible, since z-contingent transfers would change risk aversion and thus certainty-equivalent penalties for unhedgeable risk.[16]

Third, we assume normal inventory shocks and quadratic production technologies in spot markets; this drives our two-instrument "markets for Greeks" result. Under these assumptions, z-CE wealth is quadratic in spot prices. When quadratic z-CE wealth is combined with CARA utility and Gaussian spot-price risk, agents' risk-neutral distributions are Gaussian. Since a Gaussian is fully characterized by its mean and variance, futures and variance swaps align agents' entire risk-neutral distributions, delivering the price-contingent completeness result of Proposition 5. Outside our Gaussian-shock, quadratic-technology setting, matching the first two risk-neutral moments would generally not align agents' full risk-neutral distributions.[17] Our model should not be interpreted as saying futures and variance swaps generally implement optimal allocations. Rather, our model simply gives up on analyzing risk beyond second moments; the upside is that we obtain a very tractable model of risk up to second-order effects.

In summary, quasilinearity eliminates wealth effects in spot-market demand; CARA utility eliminates wealth effects in absolute risk aversion; and Gaussian shocks with quadratic technologies reduce the price-risk problem to first and second moments. The payoff is a set of analytically tractable and economically intuitive solutions for arbitrary input parameters.

Macrofinance with Idiosyncratic Risks. We loosely relate to a macrofinance literature on uninsurable idiosyncratic risk and asset prices (Constantinides and Duffie, 1996; Heaton

governed primarily by production technologies and inventory capacities. It is less natural for household-level consumption of goods that represent large expenditure shares. In this respect, our model is closer to the partial-equilibrium and industrial-organization tradition than to the macroeconomic modelling tradition.

[16] Quasilinearity and CARA utility help us avoid classic pathologies in the GEI literature (Hart, 1975; Stiglitz, 1982; Geanakoplos and Polemarchakis, 1986; Dávila and Korinek, 2018; Magill and Quinzii, 2002). Quasilinearity separates spot allocation from financial risk-sharing, reducing the post-spot problem to a one-good risk-sharing problem. Thus, equilibria with suitably bounded cash-settled assets are constrained-efficient in the sense of Geanakoplos and Polemarchakis (1986): reallocating existing financial assets cannot create a Pareto improvement. Under CARA utility, nonredundant cash-settled assets also generate strictly monotone aggregate demand, implying unique equilibrium prices and payoff allocations, up to redundant portfolio decisions.

[17] As a possible direction for future work, one could extend our model to non-Gaussian shocks, for example by imposing a finite-dimensional exponential-family structure on the induced price-risk problem, in the spirit of Breon-Drish (2015). Such an extension could potentially allow for the study of higher-order price derivatives, with payoffs designed to span the relevant higher-order sufficient statistics.

and Lucas, 1996). Krueger and Lustig (2010) demonstrate that the scale-invariance of CRRA preferences allows *multiplicative* idiosyncratic shocks to be elegantly factored out of asset prices. The *level*-invariance of CARA utility plays an analogous technical role in our analysis; *additive* idiosyncratic shocks in our model affect equilibrium outcomes, but in an analytically tractable manner captured by our z-CE wealth penalties.

The Krueger and Lustig setting is more suitable for general-equilibrium macroeconomic questions, for which income effects are central. In contrast, we think of our setting as a *partial equilibrium model of risk*: consumers are risk-averse, but income transfers affect neither spot market behavior nor consumers' absolute risk aversion over dollar wealth.

Differences from CARA-normal and Inventory-Risk Models. In the standard "CARA-normal" setting, y_i represents a *financial asset* with monetary payoff ψy_i, where ψ is random. Agents trade the asset before ψ is realized to share return risk; after realization, the asset is equivalent to money, so there is no meaningful spot market.

In our setting, y_i represents a generalized *commodity*, such as oil or wheat, which is *converted* to money-equivalents through the agent-dependent production technology in (2). Even after inventory shocks x_i are realized and uncertainty is resolved, consumers trade to improve *allocative efficiency*, moving goods toward those best able to convert them into money.

We assume there is no fundamental link between a consumer's ability to elastically supply commodities (κ_i) and their tolerance for large fluctuations in wealth (α_i^{-1}). This separation formalizes the classical narrative that derivative markets facilitate a "division of labor" between the real and financial aspects of commodity trade. A physical commodity processor may have made large physical infrastructure investments, resulting in high commodity processing capacity ($\kappa_i \uparrow$); however, due to small firm size or high leverage, the processor may have low tolerance for financial risk ($\alpha_i^{-1} \downarrow$). Conversely, a large hedge fund may have significant financial capital and risk tolerance ($\alpha_i^{-1} \uparrow$), but limited ability to trade physical commodities ($\kappa_i \downarrow$). Derivative markets improve social welfare by allowing the former group to handle physical commodity processing, while offloading the resultant wealth risks to the latter group (Bernstein, 1996; Shiller, 2012).

This separation is particularly natural in commodity markets, such as energy or agriculture, where storage is often costly or capacity-constrained, so financial derivatives are an important mechanism through which producers and consumers hedge the wealth risks generated by future spot market trade.[18]

[18] When markets are competitive and physical carry is feasible, inventory accumulation can serve as a partial substitute for futures trading. Our framework abstracts from this force entirely for simplicity. Note,

Our model also abstracts entirely from common values, information frictions, and other forces analyzed in the rational expectations equilibrium literature following Grossman and Stiglitz (1980). We assume futures are traded before agents receive any signals, so there is no scope for adverse selection. The only friction in our setting is market incompleteness: financial contracts exist, but do not span the full state space.

We also abstract from other theoretical forces studied in the literature, such as dynamics, market power, moral hazard, and imperfect rationality. Combining such elements with our framework may be an interesting direction for future work.

9 Conclusion

Modern financial markets are to a large extent markets for price risk. This paper constructs and solves a model of risk sharing in these markets. Spot market equilibria, in optimizing the allocation of goods across agents, naturally endow agents with certain risk exposures to spot prices. The role of financial markets is to allow agents to trade these exposures; we show that price risk markets achieve *price-contingent completeness*, optimally sharing the components of wealth risk that are spanned by spot prices. But price derivatives do not span the state space, and leave agents exposed to unhedgeable idiosyncratic risks that are orthogonal to spot market prices. Financial asset pricing, in equilibrium, is driven by society's aggregate exposure to price risks, adjusted for the effects of the incompleteness of price risk markets.

More broadly, our framework provides a template for applied incomplete-markets theory. The combination of CARA utility, quadratic production technologies, and normal shocks sidesteps many of the known technical difficulties in the incomplete markets literature (Hart, 1975; Geanakoplos and Polemarchakis, 1986), resulting in a model rich enough to study asset trade, pricing, and welfare, yet tractable enough to solve analytically. While this paper focuses on price derivatives, we believe our core theoretical tools – z-certainty equivalents, Gaussian risk-neutral measures, and the resulting "markets for Greeks" characterization of equilibrium – can be generalized to the analysis of other financial assets. We hope this contribution stimulates further work developing applied models of risk sharing in incomplete financial markets.

however, that in our two-period setting storage cannot replicate the role of variance swaps, which transfer second-order price risk rather than the physical commodity itself.

References

Acharya, Viral V., Lars A. Lochstoer, and Tarun Ramadorai. 2013. "Limits to arbitrage and hedging: Evidence from commodity markets." *Journal of Financial Economics*, 109(2): 441–465.

Arrow, Kenneth J. 1964. "The role of securities in the optimal allocation of risk-bearing." *The Review of Economic Studies*, 31(2): 91–96.

Baesel, Jerome, and Dwight Grant. 1982. "Equilibrium in a futures market." *Southern Economic Journal*, 49(2): 320–329.

Bernstein, Peter L. 1996. *Against the Gods: The Remarkable Story of Risk.* New York: Wiley.

Black, Fischer. 1976. "The pricing of commodity contracts." *Journal of Financial Economics*, 3(1–2): 167–179.

Black, Fischer, and Myron Scholes. 1973. "The pricing of options and corporate liabilities." *Journal of Political Economy*, 81(3): 637–654.

Borch, Karl. 1962. "Equilibrium in a reinsurance market." *Econometrica*, 30(3): 424–444.

Breeden, Douglas T. 1980. "Consumption risk in futures markets." *Journal of Finance*, 35(2): 503–520.

Brennan, Michael J, and H Henry Cao. 1996. "Information, trade, and derivative securities." *The Review of Financial Studies*, 9(1): 163–208.

Breon-Drish, Bradyn. 2015. "On existence and uniqueness of equilibrium in a class of noisy rational expectations models." *The Review of Economic Studies*, 82(3): 868–921.

Britto, Ronald. 1984. "The simultaneous determination of spot and futures prices in a simple model with production risk." *The Quarterly Journal of Economics*, 99(2): 351–365.

Chabakauri, Georgy, Kathy Yuan, and Konstantinos E Zachariadis. 2022. "Multi-asset noisy rational expectations equilibrium with contingent claims." *The Review of Economic Studies*, 89(5): 2445–2490.

Constantinides, George M, and Darrell Duffie. 1996. "Asset pricing with heterogeneous consumers." *Journal of Political Economy*, 104(2): 219–240.

Correia-da-Silva, João. 2015. "Two-period economies with price-contingent deliveries." *Economic Theory*, 59(3): 509–525.

Dávila, Eduardo, and Anton Korinek. 2018. "Pecuniary externalities in economies with financial frictions." *The Review of Economic Studies*, 85(1): 352–395.

Dierker, Egbert, and Hildegard Dierker. 2010. "Welfare and efficiency in incomplete market economies with a single firm." *Journal of Mathematical Economics*, 46(5): 652–665.

Dierker, Egbert, and Hildegard Dierker. 2012. "Ownership structure and control in incomplete market economies with transferable utility." *Economic Theory*, 51(3): 713–728.

Ekeland, Ivar, Delphine Lautier, and Bertrand Villeneuve. 2019. "Hedging pressure and speculation in commodity markets." *Economic Theory*, 68(1): 83–123.

Gârleanu, Nicolae, Lasse Heje Pedersen, and Allen M Poteshman. 2009. "Demand-based option pricing." *The Review of Financial Studies*, 22(10): 4259–4299.

Geanakoplos, John D., and Heraklis M. Polemarchakis. 1986. "Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete." In *Uncertainty, Information and Communication: Essays in Honor of Kenneth J. Arrow*. Vol. 3, ed. Walter P. Heller, Ross M. Starr and David A. Starrett, 65–96. Cambridge: Cambridge University Press.

Goldstein, Itay, and Liyan Yang. 2022. "Commodity financialization and information transmission." *Journal of Finance*, 77(5): 2613–2667.

Goldstein, Itay, Yan Li, and Liyan Yang. 2014. "Speculation and hedging in segmented markets." *The Review of Financial Studies*, 27(3): 881–922.

Gorton, Gary B., Fumio Hayashi, and K. Geert Rouwenhorst. 2013. "The fundamentals of commodity futures returns." *Review of Finance*, 17(1): 35–105.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the impossibility of informationally efficient markets." *American Economic Review*, 70(3): 393–408.

Hart, Oliver D. 1975. "On the optimality of equilibrium when the market structure is incomplete." *Journal of Economic Theory*, 11(3): 418–443.

Heal, Geoffrey. 2017. "Price uncertainty and price-contingent securities." National Bureau of Economic Research Working Paper 23723.

Heaton, John, and Deborah J Lucas. 1996. "Evaluating the effects of incomplete markets on risk sharing and asset pricing." *Journal of Political Economy*, 104(3): 443–487.

Henrotte, Philippe. 1996. "Construction of a state space for interrelated securities with an application to temporary equilibrium theory." *Economic Theory*, 8(3): 423–459.

Hirshleifer, David. 1988*a*. "Residual risk, trading costs, and commodity futures risk premia." *The Review of Financial Studies*, 1(2): 173–193.

Hirshleifer, David. 1988*b*. "Risk, futures pricing, and the organization of production in commodity markets." *Journal of Political Economy*, 96(6): 1206–1220.

Hirshleifer, David. 1989*a*. "Determinants of hedging and risk premia in commodity futures markets." *Journal of Financial and Quantitative Analysis*, 24(3): 313–331.

Hirshleifer, David. 1989*b*. "Futures trading, storage, and the division of risk: A multiperiod analysis." *The Economic Journal*, 99(397): 700–719.

Hirshleifer, David. 1990. "Hedging pressure and futures price movements in a general equilibrium model." *Econometrica*, 58(2): 411–428.

Hirshleifer, Jack. 1971. "The private and social value of information and the reward to inventive activity." *American Economic Review*, 61(4): 561–574.

Hirshleifer, Jack. 1977. "The theory of speculation under alternative regimes of markets." *Journal of Finance*, 32(4): 975–999.

Keynes, John Maynard. 1930. *A Treatise on Money.* Macmillan & Company.

Krueger, Dirk, and Hanno Lustig. 2010. "When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)?" *Journal of Economic Theory*, 145(1): 1–41.

Kurz, Mordecai, and Ho-Mou Wu. 1996. "Endogenous uncertainty in a general equilibrium model with price contingent contracts." *Economic Theory*, 8(3): 461–488.

Kyle, Albert S. 2007. "Cash settlement, price manipulation, and the Modigliani-Miller theorem." Working paper.

Lyu, Chen, Marzena J Rostek, and Ji Hee Yoon. 2023. "Futures contracts and imperfect competition." Working paper available at SSRN: 4523129.

Magill, Michael, and Martine Quinzii. 2002. *Theory of Incomplete Markets.* Vol. 1, MIT Press.

Marshall, Alfred. 1920. *Principles of Economics.* 8th ed., London: Macmillan.

Martin, Ian. 2011. "Simple variance swaps." National Bureau of Economic Research Working Paper 16884.

Martin, Ian. 2017. "What is the expected return on the market?" *The Quarterly Journal of Economics*, 132(1): 367–433.

McKinnon, Ronald I. 1967. "Futures markets, buffer stocks, and income stability for primary producers." *Journal of Political Economy*, 75(6): 844–861.

Rolfo, Jacques. 1980. "Optimal hedging under price and quantity uncertainty: The case of a cocoa producer." *Journal of Political Economy*, 88(1): 100–116.

Rostek, Marzena, and Ji Hee Yoon. 2021. "Exchange design and efficiency." *Econometrica*, 89(6): 2887–2928.

Rostek, Marzena, and Ji Hee Yoon. 2024. "Innovation in decentralized markets: technology versus synthetic products." *American Economic Journal: Microeconomics*, 16(1): 63–109.

Rostek, Marzena, and Ji Hee Yoon. 2025. "Financial product design in decentralized markets." *Journal of Political Economy*, 133(3): 888–934.

Shiller, Robert J. 2012. *Finance and the Good Society.* Princeton University Press.

Stiglitz, Joseph E. 1982. "The inefficiency of the stock market equilibrium." *The Review of Economic Studies*, 49(2): 241–261.

Svensson, Lars E. O. 1981. "Efficiency and speculation in a model with price-contingent contracts." *Econometrica*, 49(1): 131–151.

Vives, Xavier. 1987. "Small income effects: A Marshallian theory of consumer surplus and downward sloping demand." *The Review of Economic Studies*, 54(1): 87–103.

Weller, Paul, and Makoto Yano. 1987. "Forward exchange, futures trading, and spot price variability: A general equilibrium approach." *Econometrica*, 55(6): 1433–1450.

Wilson, Robert. 1968. "The theory of syndicates." *Econometrica*, 36(1): 119–132.

Zhang, Anthony Lee. 2022. "Competition and manipulation in derivative contract markets." *Journal of Financial Economics*, 144(2): 396–413.

Internet Appendix

A Omitted Proofs and Derivations

A.1 Justification of Mean-Zero Inventory Shocks

It is without loss of generality to assume (5) – that the aggregate inventory shock, $\sum_{i=1}^{N} x_i$, has mean 0 – because of a redundancy in the way we specify consumers' wealth W_i: the linear term ψ and the inventory shock x_i can be renormalized in a way that keeps consumer utility unchanged. We state this in the following simple claim, which we prove in Appendix A.1.1 below.

Claim 1. For any constant A, define:

$$\tilde{\psi} \equiv \psi + A, \;\; \tilde{x}_i \equiv x_i + \kappa_i A \tag{88}$$

Then consumer i's wealth – ignoring the price term $-pq_i$, which is unaffected – can be written as:

$$W_i = \tilde{\psi}\left(\tilde{x}_i + q_i\right) - \frac{\left(\tilde{x}_i + q_i\right)^2}{2\kappa_i} + C_i\left(A\right) \tag{89}$$

where $C_i\left(A\right)$ is a constant that does not depend on q_i or x_i.

Claim 1 implies that we can "renormalize" the constant term ψ, increasing it by any constant A across all consumers, as long as we correspondingly renormalize inventory shocks x_i. Intuitively, since ψx_i is simply a linear component of preferences, increasing ψ by A can be offset by shifting each x_i by $\kappa_i A$, up to an additive constant in wealth. Since the scaling in (88) is linear in A, this immediately implies that, for any original inventory shocks x_i whose expected sum is nonzero, we can choose A to renormalize ψ and the inventory shocks so that the resulting inventory shocks have mean zero in aggregate. This choice of A is simply:

$$\sum_{i=1}^{N} E\left[\tilde{x}_i\right] = \sum_{i=1}^{N} E\left[x_i\right] + A\sum_{i=1}^{N} \kappa_i = 0$$

$$\implies A = -\frac{\sum_{i=1}^{N} E\left[x_i\right]}{\sum_{i=1}^{N} \kappa_i}$$

As a result, it is completely without loss of generality – that is, it is simply a renormalization of agents' utility functions – to assume that the expected sum of inventory shocks across consumers is 0, as we do in (5). As we show in (18) of Section 4, (5) is a natural normalization, because it implies that μ_i is equal to the negative of i's expected trade volume in spot markets.

A.1.1 Proof of Claim 1

Note that (88) implies:

$$x_i + q_i = (\tilde{x}_i + q_i) - \kappa_i A$$

Substituting for $(x_i + q_i)$, we can write W_i as:

$$W_i = \psi\left((\tilde{x}_i + q_i) - \kappa_i A\right) - \frac{\left((\tilde{x}_i + q_i) - \kappa_i A\right)^2}{2\kappa_i}$$

$$= \psi\left(\tilde{x}_i + q_i\right) - \psi\kappa_i A - \frac{\left(\tilde{x}_i + q_i\right)^2}{2\kappa_i} + A\left(\tilde{x}_i + q_i\right) - \frac{\kappa_i A^2}{2}$$

$$= (\psi + A)\left(\tilde{x}_i + q_i\right) - \frac{\left(\tilde{x}_i + q_i\right)^2}{2\kappa_i} - \kappa_i\left(\psi A + \frac{A^2}{2}\right)$$

This is the RHS of (89), with the constant:

$$C_i(A) \equiv -\kappa_i\left(\psi A + \frac{A^2}{2}\right)$$

A.2 Derivation of Spot Market Equilibrium Wealth (19)

Copying (13), consumers' wealth is:

$$W_i = \psi\left(x_i + q_i\right) - \frac{\left(x_i + q_i\right)^2}{2\kappa_i} - pq_i \tag{90}$$

Rearranging slightly, we have:

$$W_i = \psi x_i - \frac{\left(x_i + q_i\right)^2}{2\kappa_i} - (p - \psi)\, q_i$$

Substituting equilibrium quantities (17) and prices (16), we have:

$$W_i = \psi x_i - \frac{\left(\frac{\kappa_i}{\sum_{j=1}^N \kappa_j}\sum_{j=1}^N x_j\right)^2}{2\kappa_i} - \left(-\frac{\sum_{j=1}^N x_j}{\sum_{j=1}^N \kappa_j}\right)\left(\frac{\kappa_i}{\sum_{j=1}^N \kappa_j}\sum_{j=1}^N x_j - x_i\right)$$

Simplifying, we obtain (19).

A.3 Proof of Proposition 1

After the realization of shocks, consumers' utility is quasilinear in money, implying that all Pareto-efficient outcomes must maximize the sum of consumers' monetary-equivalent values of goods; any non-maximizing allocation is Pareto-dominated with transfers. That is, efficient allocations must solve:

$$\max_{y_i} \sum_{i=1}^{N} W_i(y_i) = \max_{y_i} \sum_{i=1}^{N} \left[\psi y_i - \frac{y_i^2}{2\kappa_i}\right]$$

$$s.t. \sum_{i=1}^{N} y_i = \sum_{i=1}^{N} x_i$$

The Lagrangian is:

$$\Lambda = \sum_{i=1}^{N} \left[\psi y_i - \frac{y_i^2}{2\kappa_i}\right] - \lambda \left(\sum_{i=1}^{N} y_i - \sum_{i=1}^{N} x_i\right)$$

The first-order condition is:

$$0 = \frac{\partial \Lambda}{\partial y_i} = \psi - \frac{y_i}{\kappa_i} - \lambda$$

This implies simply that consumers' marginal rates of substitution between wealth and goods, $\frac{\partial W_i}{\partial y_i} = \psi - \frac{y_i}{\kappa_i}$, must be equalized:

$$\frac{y_i}{\kappa_i} = \psi - \lambda$$

Combining this with the resource constraint (6), we get (10), which uniquely characterizes the allocations $y_i^*(\boldsymbol{x})$ which maximize aggregate wealth, conditional on any inventory shock realization $\boldsymbol{x}$. Plugging (10) into consumers' production technology (3) and summing, we then get (11).

To show that (10) is necessary for Pareto efficiency, suppose $y_i(\boldsymbol{x})$ does not satisfy (10) for some i and $\boldsymbol{x}$. For any realization $\boldsymbol{x}$ where (10) is violated, replacing $\boldsymbol{y}(\boldsymbol{x})$ with $\boldsymbol{y}^*(\boldsymbol{x})$ increases total social wealth $W(\boldsymbol{y}(\boldsymbol{x}))$ and loosens the constraint (8). We can thus increase $G_i(\boldsymbol{x})$ for all i, leading to a Pareto improvement. We can also conclude from the above analysis that (8) must be binding.

Suppose now that $y_i(\boldsymbol{x})$ does satisfy (10). The optimal risk-sharing condition is simply the result of Borch (1962) in our setting. Pareto efficiency requires agents to equate the ratio of their marginal utilities across all states:

$$\frac{w_j}{w_i} = \frac{U_i'(G_i(\boldsymbol{x}))}{U_j'(G_j(\boldsymbol{x}))} = \frac{\alpha_i e^{-\alpha_i G_i(\boldsymbol{x})}}{\alpha_j e^{-\alpha_j G_j(\boldsymbol{x})}},$$

where w_i is the weight for i's utility. Combining this with the binding constraint (8), we get (12), following Wilson (1968). Hence, the wealth allocation is uniquely characterized up to agent-specific, state-independent constants.

Note that an additional technical integrability condition imposed on primitives is required for expected utility in the first-best outcome to be finite; we state this condition as (84) of Proposition 6.

A.4 Derivation of Projection Coefficients in Definition 1

Expression (16) for prices gives us:

$$Var\left(z\right)=\frac{\sum_i \sigma_i^2}{\left(\sum_i \kappa_i\right)^2}, Cov\left(\epsilon_i, z\right)=-\frac{\sigma_i^2}{\sum_j \kappa_j}$$

The coefficient β_i in (31) is just the OLS regression coefficient of ϵ_i on z:

$$\beta_i=\frac{Cov\left(\epsilon_i, z\right)}{Var\left(z\right)}=-\left(\sum_j \kappa_j\right)\frac{\sigma_i^2}{\sum_j \sigma_j^2}$$

The residual variance can be calculated by subtracting the variance of $\beta_i z$ from the variance of ϵ_i:

$$\nu_i \equiv Var\left(\epsilon_i \mid z\right)=Var\left(\epsilon_i\right)-\frac{\left(Cov\left(\epsilon_i, z\right)\right)^2}{Var\left(z\right)}$$

$$\nu_i=\sigma_i^2\left(\frac{\sum_{j\neq i}\sigma_j^2}{\sum_j \sigma_j^2}\right)$$

A.5 $V\left(q^c, z^c\right)$ Must Be Positive

Consider a portfolio which is long 1 unit of the variance swap and short $2z^c$ units of the futures contract. The payoff of this portfolio is:

$$\underbrace{\left(z^2-q^c\right)}_{Variance\ Swap}-\underbrace{\left(2z^c z-2\left(z^c\right)^2\right)}_{Futures}$$

$$=\underbrace{\left(z-z^c\right)^2}_{Payoff}-\underbrace{\left(q^c-\left(z^c\right)^2\right)}_{Price} \tag{91}$$

The random payoff $\left(z-z^c\right)^2$ of this position is positive with probability 1. Thus, the constant price $q^c-\left(z^c\right)^2$, which we defined as $V\left(q^c, z^c\right)$, must also be positive; if $q^c \leq \left(z^c\right)^2$, then this portfolio is an arbitrage – it has positive payoffs net of its price with probability 1 – and demand for the portfolio is unbounded.

A.6 Proof of Proposition 4

With two derivatives, (47) must hold for both simultaneously, implying:

$$E_{\mathbb{Q}_i^{c_i,d_i}}\left(z\right) = z^c, E_{\mathbb{Q}_i^{c_i,d_i}}\left(z^2\right) = q^c$$

Now, the $\mathbb{Q}_i^{c_i,d_i}$-expectation of the ψ-centered "variance swap", z^2, can be written as:

$$E_{\mathbb{Q}_i^{c_i,d_i}}\left(z^2\right) = \left(E_{\mathbb{Q}_i^{c_i,d_i}}\left(z\right)\right)^2 + Var_{\mathbb{Q}_i^{c_i,d_i}}\left(z\right) = \left(z^c\right)^2 + Var_{\mathbb{Q}_i^{c_i,d_i}}\left(z\right) \tag{92}$$

where the $(z^c)^2$ correction term arises because z has $\mathbb{P}$-mean 0, but not $\mathbb{Q}_i^{c_i,d_i}$-mean 0. Substituting for $Var_{\mathbb{Q}_i^{c_i,d_i}}(z)$ in (92) using (69) and rearranging, we have:

$$\frac{1}{\frac{1}{\sigma_p^2} + 2\alpha_i\left(\Gamma_i + d_i\right)} = q^c - \left(z^c\right)^2 \tag{93}$$

That is, the $\mathbb{Q}_i^{c_i,d_i}$-implied variance on the LHS must equal the q^c, z^c-implied price of variance on the RHS, which we define as $V\left(q^c, z^c\right)$ in (71). Solving for d_i, we obtain (72). Note that, as long as the $\mathbb{Q}_i^{c_i,d_i}$-variance $V\left(q^c, z^c\right)$ is positive, we have:

$$d_i\left(q^c, z^c\right) + \Gamma_i = \frac{1}{2\alpha_i}\left(\frac{1}{V\left(q^c, z^c\right)} - \frac{1}{\sigma_p^2}\right) > -\frac{1}{2\alpha_i\sigma_p^2}$$

Hence the regularity condition (67) in Lemma 4 holds. Intuitively, i purchases contracts until her $\mathbb{Q}$-variance, which is (69), is equal to the market "price of variance" $V\left(q^c, z^c\right)$; i's optimized $\mathbb{Q}$-variance is thus finite whenever the market price is positive. While certain choices of d_i deliver unboundedly negative utility to i, i would never optimally purchase these values of d_i. The assumption that $V\left(q^c, z^c\right)$ is positive is needed for the consumer's optimization problem to be well-posed; we show in Appendix A.5 that, if $V\left(q^c, z^c\right)$ is negative, an arbitrage opportunity exists and contract demand is thus unbounded.

The derivation of linear contract demand is analogous to Proposition 3. The $\mathbb{Q}_i^{c_i,d_i}$-mean of z is:

$$-\alpha_i\left(\Delta_i + c_i\right)V_i$$

with V_i defined in (69) of Lemma 4 as the $\mathbb{Q}_i^{c_i,d_i}$-variance of z. But (93) implies that this

$\mathbb{Q}_i^{c_i,d_i}$-variance is just equal to V as defined in (71); hence, futures demand with variance swaps, (73), is simply futures demand without variance swaps, (56) of Proposition 3, with V in place of V_i.

Imposing market clearing on variance swap demand in (72), we find that the equilibrium price of variance must satisfy:

$$\left(\frac{1}{V\left(q^c_{eq}, z^c_{eq}\right)} - \frac{1}{\sigma_p^2}\right)\sum_i \frac{1}{2\alpha_i} = \sum_i \Gamma_i \tag{94}$$

Intuitively, the variance risk premium $\frac{1}{V} - \frac{1}{\sigma_p^2}$ depends on the aggregate gamma position $\sum_i \Gamma_i$, and risk aversion $\sum_i \frac{1}{2\alpha_i}$. Substituting for Γ_i using expression (41), we have:

$$\left(\frac{1}{V\left(q^c_{eq}, z^c_{eq}\right)} - \frac{1}{\sigma_p^2}\right)\sum_i \frac{1}{2\alpha_i} = \sum_i \left(\beta_i + \frac{\kappa_i}{2} - \frac{\alpha_i}{2}\nu_i\right)$$

Applying (33) and simplifying, we obtain (75). Rearranging, a necessary and sufficient condition for $V\left(q^c_{eq}, z^c_{eq}\right)$ to be positive – so consumer demand in the candidate equilibrium is well-posed – is (74). Intuitively, this condition requires aggregate risk aversion $\left(\sum_i \frac{1}{\alpha_i}\right)^{-1}$ to be low relative to the variance of prices σ_p^2, and an "aggregate concavity" term, $\sum_i (\kappa_i + \alpha_i\nu_i)$, which depends on the sum of inventory capacities κ_i and basis risks $\alpha_i\nu_i$.[19]

The linear contract price (76) follows from an argument analogous to the proof of Proposition 3.

Expressions (75) and (76) show that z^c-centered variance V turns out to be analytically simpler than the ψ-centered variance q^c, because z^c is the endogenous common mean of consumers' risk-neutral distributions; this is why we state Proposition 4 in terms of $V(q^c, z^c)$. To solve for equilibrium prices q^c_{eq}, z^c_{eq}, we would calculate the equilibrium risk-neutral variance $V\left(q^c_{eq}, z^c_{eq}\right)$ using (75), calculate the risk-neutral mean z^c_{eq} using (76), and then solve for the implied ψ-centered swap price q^c_{eq} through the $\mathbb{Q}$-variance identity (71).

[19]Note that (74) can be obtained by summing the individual integrability conditions (51) in Lemma 3 over i and simplifying using (41) and (33). Thus, requiring (51) for all i is sufficient, but not necessary, for an equilibrium with variance swaps to exist. Intuitively, one can construct cases in which an agent's expected utility is $-\infty$ in the absence of variance swap markets, because her pre-trade variance exposure violates (51); yet variance swaps allow her to offload enough variance risk to other agents that all agents' post-trade variance exposures satisfy the integrability condition, and hence a variance swap equilibrium exists.

A.7 Proof of Lemma 5

Rearranging (72), optimal variance swap purchasing implies:

$$\frac{1}{\alpha_i V} = \frac{1}{\alpha_i \sigma_p^2} + 2\left(\Gamma_i + d_i\right) \tag{95}$$

Plugging (95) into futures demand (73) and summing, we can write the futures market clearing condition as:

$$0 = \sum_{i=1}^{N} c_i\left(q^c, z^c\right) = \sum_{i=1}^{N} \left[-\mu_i + \psi\left(\nu_i \alpha_i - \beta_i\right) - z^c \left(\frac{1}{\alpha_i \sigma_p^2} + 2\left(\Gamma_i + d_i\right)\right)\right] \tag{96}$$

Now, variance swap market clearing requires $\sum_i d_i = 0$, so variance swap positions simply drop out of (96), leaving us with the market clearing condition:

$$\sum_{i=1}^{N} \left[-\mu_i + \psi\left(\nu_i \alpha_i - \beta_i\right) - z^c \left(\frac{1}{\alpha_i \sigma_p^2} + 2\Gamma_i\right)\right] = 0$$

which is exactly the market clearing condition for futures contracts without variance swaps, (60) in the proof of Proposition 3.

A.8 Proof of Lemma 6

Optimal risk-sharing is characterized by the Borch (1962) rule, which is (12) of Proposition 1. Substituting into (81), letting W denote random total social wealth to be shared, we have:

$$\begin{aligned} CE_i &= -\frac{1}{\alpha_i} \log\left[E\left[\exp\left[-\alpha_i\left(C_i + \frac{\alpha_i^{-1}}{\sum_j \alpha_j^{-1}} W\right)\right]\right]\right] \\ &= C_i - \frac{1}{\alpha_i} \log E\left[\exp\left(-\frac{1}{\sum_j \alpha_j^{-1}} W\right)\right] \end{aligned}$$

Hence,

$$\sum_{i=1}^{N} CE_i = \sum_{i=1}^{N} C_i - \left(\sum_{i=1}^{N} \alpha_i^{-1}\right) \log E\left[\exp\left(-\frac{1}{\sum_j \alpha_j^{-1}} W\right)\right]$$

Now from Proposition 1, we assumed the lump-sum transfers satisfy $\sum_{i=1}^{N} C_i = 0$, hence we recover (82).

A.9 Proof of Proposition 6

It is analytically convenient to work with the non-normalized spot price $p \equiv \psi + z$, so we will use this throughout most of this proof.

A.9.1 First-Best CE

Substituting the spot price (27) into first-best wealth (11), we have:

$$W^*(p) = \left(\psi \left(\sum_{i=1}^{N} \kappa_i\right)(\psi - p)\right) - \frac{\left(\sum_{i=1}^{N} \kappa_i\right)^2 (\psi - p)^2}{2\sum_{i=1}^{N} \kappa_i}$$

$$W^*(p) = \frac{\sum_{i=1}^{N} \kappa_i}{2}\left(\psi^2 - p^2\right) \tag{97}$$

Hence, from Lemma 6, we have:

$$CE^{FB} = -\frac{1}{\alpha^{agg}} \log E\left[\exp\left[-\alpha^{agg}\left(\frac{\sum_{i=1}^{N} \kappa_i}{2}\left(\psi^2 - p^2\right)\right)\right]\right]$$

$$CE^{FB} = \frac{\sum_{i=1}^{N} \kappa_i}{2}\psi^2 - \frac{1}{\alpha^{agg}} \log\left(E\left[\exp\left[\alpha^{agg}\left(\frac{\sum_{i=1}^{N} \kappa_i}{2} p^2\right)\right]\right]\right) \tag{98}$$

Now, $p \sim N\left(\psi, \sigma_p^2\right)$. In general, we have:

$$E\left[e^{tp^2}\right] = \frac{1}{\sqrt{1 - 2t\sigma_p^2}} \exp\left(\frac{t\psi^2}{1 - 2t\sigma_p^2}\right) \tag{99}$$

and the expectation is finite whenever:

$$2t\sigma_p^2 < 1 \tag{100}$$

Hence,

$$E\left[\exp\left[\alpha^{agg}\left(\frac{\sum_{i=1}^{N} \kappa_i}{2} p^2\right)\right]\right] = \frac{1}{\sqrt{1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^{N} \kappa_i}} \exp\left(\frac{\alpha^{agg}\psi^2 \sum_{i=1}^{N} \kappa_i}{2\left(1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^{N} \kappa_i\right)}\right)$$

Plugging into (98) and simplifying, we obtain (85). The integrability condition (84) corresponds to (100).

A.9.2 Derivatives Equilibrium CE

First, we note that z-CE wealths can be thought of as the conditional expectation of W_i^{Spot} given z, minus a variance penalty term. Formally, notice that, since we defined η_i so that $E[\eta_i \mid z] = 0$, we have from (34):

$$E\left[W_i^{Spot} \mid z\right] = (\psi + z)(\mu_i + \beta_i z) + \frac{\kappa_i}{2} z^2$$

Hence, from Definition 2, we can write:

$$\bar{W}_i = E\left[W_i^{Spot} \mid z\right] - \frac{\alpha_i}{2} \nu_i (\psi + z)^2$$

Now, we wish to calculate aggregate z-CE wealth as a function of z, and compare it to first-best wealth $W^*(z)$. Note that:

$$W^*(z) = E[W^*(z) \mid z] = E\left[\sum_{i=1}^{N} W_i^{Spot} \mid z\right] = \sum_{i=1}^{N} E\left[W_i^{Spot} \mid z\right] \tag{101}$$

where the last equality follows from linearity of expectations. Thus,

$$\sum_{i=1}^{N} \bar{W}_i = \sum_{i=1}^{N} \left[E\left[W_i^{Spot} \mid z\right] - \frac{\alpha_i}{2} \nu_i (\psi + z)^2\right] = W^*(z) - (\psi + z)^2 \sum_{i=1}^{N} \frac{\alpha_i \nu_i}{2} \tag{102}$$

Intuitively, we have shown that $\sum_{i=1}^{N} \bar{W}_i$, the sum of z-CEs across agents, is equal to first-best wealth $W^*(z)$, minus an "unspanned risk" cost equal to the sum of basis risk penalties across agents; this is the squared non-normalized spot price $(\psi + z)^2$, times half the sum of $\alpha_i \nu_i$ across agents. Using (97), and noting that $p = \psi + z$, we can further write:

$$\sum_{i=1}^{N} \bar{W}_i = \frac{\sum_{i=1}^{N} \kappa_i}{2} \psi^2 - \frac{\sum_{i=1}^{N} (\kappa_i + \alpha_i \nu_i)}{2} p^2 \tag{103}$$

We have thus shown that society's aggregate z-CE wealth takes the form (103). Proposition 5 shows that, when futures and variance swaps are traded, the Borch (1962) rule holds and z-CE wealth is optimally shared across agents. Lemma 6 thus holds, and the N agents behave like a representative agent with risk aversion α^{agg} receiving random wealth (103). We can thus calculate the representative agent's certainty equivalent analogously to above:

$$CE^{Equil} = -\frac{1}{\alpha^{agg}} \log E\left[\exp\left[-\alpha^{agg}\left(\frac{\sum_{i=1}^{N} \kappa_i}{2} \psi^2 - \frac{\sum_{i=1}^{N} (\kappa_i + \alpha_i \nu_i)}{2} p^2\right)\right]\right]$$

Using the Gaussian integral formula (99) and simplifying, we get (86).

A.9.3 Approximating the Difference

Subtracting (86) from (85), we have:

$$
\begin{aligned}
&CE^{FB} - CE^{Equil} = \\
&\qquad \frac{1}{2\alpha^{agg}} \log \left(\frac{1 - \alpha^{agg} \sigma_p^2 \sum_{i=1}^N \kappa_i}{1 - \alpha^{agg} \sigma_p^2 \sum_{i=1}^N (\kappa_i + \alpha_i \nu_i)} \right) - \\
&\qquad \left(\frac{\psi^2 \sum_{i=1}^N \kappa_i}{2 \left(1 - \alpha^{agg} \sigma_p^2 \sum_{i=1}^N \kappa_i\right)} - \frac{\psi^2 \sum_{i=1}^N (\kappa_i + \alpha_i \nu_i)}{2 \left(1 - \alpha^{agg} \sigma_p^2 \sum_{i=1}^N (\kappa_i + \alpha_i \nu_i)\right)} \right)
\end{aligned}
$$

The Taylor expansion of the logarithm is:

$$\log (1 - \epsilon) = -\epsilon + O\left(\epsilon^2\right)$$

Hence, we have:

$$
\begin{aligned}
&\frac{1}{2\alpha^{agg}} \log \left(\frac{1 - \alpha^{agg} \sigma_p^2 \sum_{i=1}^N \kappa_i}{1 - \alpha^{agg} \sigma_p^2 \sum_{i=1}^N (\kappa_i + \alpha_i \nu_i)} \right) = \\
&\qquad \frac{1}{2\alpha^{agg}} \left(- \left(\alpha^{agg} \sigma_p^2 \sum_{i=1}^N \kappa_i \right) - \left(-\alpha^{agg} \sigma_p^2 \sum_{i=1}^N (\kappa_i + \alpha_i \nu_i) \right) \right) + O\left(\epsilon^2\right) \\
&\qquad = \frac{\sigma_p^2 \sum_{i=1}^N \alpha_i \nu_i}{2} + O\left(\epsilon^2\right) \qquad (104)
\end{aligned}
$$

For the second term, we use the Taylor expansion of the geometric series:

$$(1 - \epsilon)^{-1} = 1 + \epsilon + O\left(\epsilon^2\right)$$

which gives:

$$\frac{\psi^2 \sum_{i=1}^N \kappa_i}{2 \left(1 - \alpha^{agg} \sigma_p^2 \sum_{i=1}^N \kappa_i\right)} = \left(\frac{\psi^2 \sum_{i=1}^N \kappa_i}{2} \right) \left(1 + \epsilon + O\left(\epsilon^2\right)\right)$$

Since ψ^2 itself is $O(\epsilon)$ by assumption, we then have:

$$\left(\frac{\psi^2 \sum_{i=1}^N \kappa_i}{2} \right) \left(1 + \epsilon + O\left(\epsilon^2\right)\right) = \left(\frac{\psi^2 \sum_{i=1}^N \kappa_i}{2} \right) + O\left(\epsilon^2\right)$$

Intuitively, this allows us to simply ignore the denominator in the second term, since it introduces a correction of order $O\left(\epsilon^2\right)$. We then have:

$$\frac{\psi^2 \sum_{i=1}^N \left(\kappa_i + \alpha_i \nu_i\right)}{2\left(1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^N \left(\kappa_i + \alpha_i \nu_i\right)\right)} - \frac{\psi^2 \sum_{i=1}^N \kappa_i}{2\left(1 - \alpha^{agg}\sigma_p^2 \sum_{i=1}^N \kappa_i\right)} =$$

$$\frac{\psi^2}{2}\left(\left[\sum_{i=1}^N \left(\kappa_i + \alpha_i \nu_i\right)\right] - \sum_{i=1}^N \kappa_i\right) + O\left(\epsilon^2\right) = \frac{\psi^2}{2}\sum_{i=1}^N \alpha_i \nu_i + O\left(\epsilon^2\right) \quad (105)$$

Combining (104) and (105), we obtain:

$$CE^{FB} - CE^{Equil} = \frac{\left(\sigma_p^2 + \psi^2\right)\sum_{i=1}^N \alpha_i \nu_i}{2} + O\left(\epsilon^2\right)$$

Since $E\left[p^2\right] = \sigma_p^2 + \psi^2$, we have derived (87).

A.10 Welfare in Isolated Spot Markets

From Lemma 1 and Definition 2, the exponential moment used to compute consumer i's expected utility from her spot market wealth is:

$$E\left[\exp\left(-\alpha_i W_i^{Spot}\right)\right] = E\left[\exp\left(-\alpha_i \bar{W}_i\right)\right] = E\left[\exp\left[-\alpha_i\left(A_i + \Delta_i z + \Gamma_i z^2\right)\right]\right]$$

Now, we have $z \sim N\left(0, \sigma_p^2\right)$; thus, applying the Gaussian integral formula (99), we obtain:

$$E\left[\exp\left[-\alpha_i\left(A_i + \Delta_i z + \Gamma_i z^2\right)\right]\right] = \frac{1}{\sqrt{1 + 2\alpha_i \sigma_p^2 \Gamma_i}} \exp\left(-\alpha_i A_i + \frac{\alpha_i^2 \Delta_i^2 \sigma_p^2}{2\left(1 + 2\alpha_i \sigma_p^2 \Gamma_i\right)}\right)$$

Plugging this into the CE formula (81), i's CE for her spot market wealth is:

$$CE_i^{Spot} = A_i + \frac{1}{2\alpha_i}\log\left(1 + 2\alpha_i \sigma_p^2 \Gamma_i\right) - \frac{\alpha_i \Delta_i^2 \sigma_p^2}{2\left(1 + 2\alpha_i \sigma_p^2 \Gamma_i\right)}$$