Welfare-Improving Price Controls in Incomplete Markets*

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

February 2026

PRELIMINARY AND INCOMPLETE

This paper builds on another paper by the same authors, Han, Hu and Zhang (2026). The models in the two papers are identical, so we have copied the model exposition and some results in this paper for expositional completeness.

Abstract

Spot markets facilitate allocative efficiency, reallocating goods to maximize society's money-equivalent wealth. Financial markets facilitate risk-sharing, redistributing optimized wealth according to consumers' risk preferences. When financial markets are incomplete, the wealth generated by efficient spot markets is not optimally shared across agents. Market incompleteness implies that commodity price controls can be Pareto-improving, if they sufficiently improve risk sharing across agents.

*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

This paper analyzes the possibility of using price controls to improve social welfare through improved risk sharing in incomplete markets. We build a model in which money and commodities are traded in *spot markets*, which achieve *allocative efficiency* by redistributing the commodity across agents to maximize its money-equivalent value. Spot markets allocate commodities efficiently, but do not generally distribute wealth efficiently across uncertain states of the world. The role of idealized *financial markets* is to transform the wealth generated by efficient spot markets into state-contingent payoffs that optimally share risks across agents.

When financial markets are absent or incomplete, market outcomes are allocatively efficient, but risks generated by spot markets are not optimally shared across agents. In such settings, interventions in spot markets that reduce allocative efficiency can be welfare-enhancing, if they sufficiently improve risk sharing. We then show that *price controls* in commodity markets can be Pareto-improving, in expected-utility terms, through their benefits for risk-sharing.

Our model has two goods: "money", and a real commodity, such as oil, wheat, or steel. Agents have a concave *production technology* which converts the commodity into money-equivalents: in other words, conditional on shock realization, agents' utility is quasilinear over the commodity and money. Ex-ante, agents are risk-averse over money, with CARA utility with potentially different risk aversions. The only source of uncertainty in the economy is agents' random endowments of the commodity. This setting is very general, but can be thought of as modelling trade and monetary risk sharing in any real factor of production.

What is the social first-best outcome? The social planner, in this setting, essentially solves a two-stage problem. Conditional on any realization of shocks, market outcomes should be *allocatively efficient*: commodities should be in the hands of those agents who have the most efficient technologies to convert them into money-equivalent consumption, in order to maximize the consumption available to society in each state of the world. *Risk sharing* should be optimal: shocks to aggregate inventory, optimally filtered through agents' production technologies, imply that society faces risky aggregate money-equivalent consumption; aggregate consumption shocks should be divided proportionally across agents according to their risk aversions, following Borch (1962) and Wilson (1968).

There is a simple and classical Arrow-Debreu implementation of the first-best outcome, which can be thought of as a backward induction process. Allocative efficiency is achieved through *spot markets*, markets which open after inventory shocks are realized, in which money and commodities are traded for each other. Agents' preferences over goods are quasilinear, so Walrasian equilibrium in spot markets is unique and maximizes society's aggregated money-metric utility, as in the social planner's problem. Risk sharing is achieved through *financial markets*: when agents can trade contingent claims on states of the world – the entire vector of inventory shocks – then financial markets decentralize the social planner's first-best solution.

The core departure point of our paper, as in the classic literature on incomplete markets, is that financial markets are likely incomplete in practice, so the planner's first-best is likely unattainable. When there are no financial markets, spot markets are allocatively efficient, but the distributions of wealth induced by Walrasian equilibria in spot markets fail to efficiently distribute risk among agents. A simple way to see why this must be the case is that spot market equilibria are functions only of realized inventory shocks and production technologies; they do not depend on risk aversions, and thus spot market equilibria cannot possibly share monetary risk according to risk aversions.

The inefficiency of spot markets in our model implies that *price controls* can be Pareto-improving. Price controls unambiguously decrease allocative efficiency, as they induce rationing and deadweight loss. However, they can improve risk-sharing, since they redistribute wealth towards agents with extreme inventory shocks, who have high marginal utility of wealth. When these insurance benefits are larger than the allocative efficiency losses, all agents achieve higher expected utility under price controls relative to free spot markets.

This current paper builds on the model of Han, Hu and Zhang (2026). Sections 2 and 3, and parts of Section 4 are identical to Han, Hu and Zhang (2026); we include them here so the current paper is self-contained. The substantively new results in this paper are in Sections 5 and 6.

2 Model

The model in this paper is identical to that of Han, Hu and Zhang (2026); we repeat the model exposition unchanged here for completeness. Notationally, we will use bold symbols for vectors, for example writing $\boldsymbol{x}$ to mean the vector $(x_1 \dots x_N)$.

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

[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.

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. i is endowed with an initial constant amount m_i of money, and all consumers can hold infinitely 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:

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

CARA utility implies that money endowments m_i have no effect on i's behavior, since m_i simply scales $U_i\left(W_i\right)$ 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:

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

Wealth $W_i\left(y_i\right)$ 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 $W_i\left(y_i\right)$ 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 $W_i\left(y_i\right)$ as a preference function for y_i rather than a production technology; in these terms, consumer i gets utility equivalent to having $W_i\left(y_i\right)$ 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:

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

We will impose the normalization that the mean of the average 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 inventory shocks 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 two financial market structures: complete financial markets with Arrow securities; and no financial markets. 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 securities trades in period 1 and inventory shock realizations $x_1 \ldots x_N$, consumers trade money for the commodity.

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 $y_1\left(\boldsymbol{x}\right)$ to $y_N\left(\boldsymbol{x}\right)$, satisfying, pointwise in $\boldsymbol{x}$, the aggregate resource constraint:

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

It is of course equivalent to assume that the social planner chooses net trade amounts $q_i\left(\boldsymbol{x}\right)$ rather than final inventories $y_i\left(\boldsymbol{x}\right)$. 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:

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

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

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

Thus, in sum, the social planner chooses commodity allocations $y_i\left(\boldsymbol{x}\right)$ and money allocations $G_i\left(\boldsymbol{x}\right)$, 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, $\tilde{G}_i\left(\boldsymbol{x}\right)$ Pareto-dominates $G_i\left(\boldsymbol{x}\right)$ if:

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

To handle probability-zero edge cases, we will additionally strengthen this definition by saying that $\tilde{G}_i\left(\boldsymbol{x}\right)$ Pareto-dominates $G_i\left(\boldsymbol{x}\right)$ if $\tilde{G}_i\left(\boldsymbol{x}\right) \geq G_i\left(\boldsymbol{x}\right)$ for all i and all realizations of $\boldsymbol{x}$, and $\tilde{G}_i\left(\boldsymbol{x}\right) > G_i\left(\boldsymbol{x}\right)$ for some i and $\boldsymbol{x}$, even if the set of $\boldsymbol{x}$ values on which the inequality is strict has measure zero.

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

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

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

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

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

Optimal risk-sharing requires that wealth is shared 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, in the sense that commodities are distributed in a way which optimally produces money, given consumers' heterogeneous production technologies $W_i(y_i)$. If this were not the case for any realization $\boldsymbol{x}$, society could simply reallocate goods to generate more wealth, and redistribute this wealth to increase all consumers' money-metric utility 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; higher-κ_i consumers have more elastic demand, suffering lower costs for absorbing inventory, and thus absorb a larger fraction of aggregate inventory shocks in equilibrium.

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 spot market optimal 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 the equalization of the ratio of marginal utility across states. Under our assumption of CARA utility, the classic results of Borch (1962) and Wilson (1968) imply that any Pareto-efficient allocations redistribute risks in wealth, driven by uncertainty in $\boldsymbol{x}$, affinely according to consumers' risk aversions, as in (12).

Proposition 1 shows that Pareto efficiency is a very restrictive criterion in our model: consumers' spot market outcomes are fully pinned down, and wealths are pinned down across states up to consumer-specific constants. Thus, with slight abuse of terminology, we will occasionally refer to the outcomes described in Proposition 1 as "the first-best outcome" in singular form, implicitly 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:

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

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 attained through financial asset trade in the first period. Differentiating (13) with respect to q_i, we have:

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

(14) depends on x_i and q_i, but not C_i: preferences are quasilinear, so there are no income effects. Thus, arbitrarily wealth transfers in financial markets have no effect on spot market demand, and thus 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, in terms of money:

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

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:

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

The spot market clearing price is thus simply a function of consumers' inventory shocks:

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

Intuitively, the equilibrium price deviation from ψ is simply the aggregate inventory shock

$\sum_{i=1}^{N} x_i$, divided by the aggregate "inventory capacity", or alternatively the slope of aggregate demand, $\sum_{i=1}^{N} \kappa_i$. Plugging (16) into consumer demand (15), we can calculate consumers' equilibrium inventories:

$$x_i + q_i^{Spot}(x_1 \ldots x_N) = \frac{\kappa_i}{\sum_{j=1}^{N} \kappa_j} \sum_{j=1}^{N} x_j \tag{17}$$

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

Let W_i^0 represent i's welfare in autarky, from consuming her endowment x_i:

$$W_i^0 = C_i + \psi x_i - \frac{x_i^2}{2\kappa_i} \tag{18}$$

Taking the difference between (13) and (18), plugging in (15), and simplifying, i's money-metric welfare gains from trade are simply:

$$W_i - W_i^0 = \frac{q_i^2}{2\kappa_i} \tag{19}$$

Since preferences are quadratic, expression (19) is just i's *consumer surplus triangle*: it is half the product of her trade quantity, q_i, and her marginal WTP for the good when trading nothing, $\frac{q_i}{\kappa_i}$. Quasilinear preferences in second-stage markets imply that compensating and equivalent variation are equal to each other, and to the integral of Marshallian demand over prices, which is (19). Society's total monetary welfare gains from trade are simply the sum of surplus triangles over all consumers.

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:

$$E\left[q_i^{Spot}(\boldsymbol{x})\right] = -E[x_i] + \frac{\kappa_i}{\sum_{j=1} \kappa_j} \sum_{j=1}^{N} E[x_j] = -\mu_i \tag{20}$$

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

$$W_i^{Spot} = \psi x_i + \frac{\kappa_i}{\sum_{j=1}^N \kappa_j} \frac{\left(\sum_{j=1}^N x_j\right)^2}{2\sum_{j=1}^N \kappa_j} - x_i \frac{\sum_{j=1}^N x_j}{\sum_{j=1}^N \kappa_j} \tag{21}$$

where we have omitted the constant money term, C_i, for convenience. Alternatively, substituting (16) for $p^{Spot}(\boldsymbol{x})$ into (21) and rearranging, we have:

$$W_i^{Spot} = p^{Spot}(\boldsymbol{x})\, x_i + \frac{\kappa_i}{2}\left(p^{Spot}(\boldsymbol{x}) - \psi\right)^2 \tag{22}$$

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*, denominated in units of wealth, which fully span the state space. Markets are then complete, the welfare theorems hold, and market equilibrium implements the first-best outcome.

Working with Arrow securities in high-dimensional state-spaces is somewhat technically involved; while we will carefully define these objects here, note that we will not work with Arrow securities outside this subsection. Since our state space $\boldsymbol{x}$ is continuous, Arrow security prices constitute a *state price density* (Duffie, 2010, ch. 2), which we will refer to as $\pi(\boldsymbol{x})$.[2] Let $\theta_i(\boldsymbol{x})$ denote the *security demand function* of consumer i; that is, consumer i purchases securities paying her a net amount $\theta_i(\boldsymbol{x})$ in state $\boldsymbol{x}$. Since there is no first-stage consumption, agents trade money across states of the world by buying Arrow securities in some states and selling them in other states. i's budget constraint is that her total expenditures must integrate to 0 across states:

$$\int \pi(\boldsymbol{x})\, \theta_i(\boldsymbol{x})\, d\boldsymbol{x} = 0 \;\forall i \tag{23}$$

Note that, following tradition in the literature, we absorb the physical probability density $f(\boldsymbol{x})$ into the definition of the state price density $\pi(\boldsymbol{x})$.

Agents purchase Arrow securities to maximize expected utility subject to (23). We will require *market clearing* pointwise in $\boldsymbol{x}$; since Arrow securities are financial assets in zero net

[2] A subtle difference between our model and the canonical setting is that, since we assume there is no first-period consumption, there is no natural numeraire in our setting. Thus, we leave $\pi(\boldsymbol{x})$ defined only up to scale. An equivalent alternative approach would be to choose some arbitrary value of $\boldsymbol{x}$ as the numeraire good.

supply, asset demands must sum to zero across agents:

$$\sum_{i=1}^{N} \theta_i(\boldsymbol{x}) = 0 \;\forall \boldsymbol{x} \tag{24}$$

Equilibrium is described by a state price density $\pi(\boldsymbol{x})$ and security demands $\boldsymbol{\theta}(\boldsymbol{x})$, such that all consumers are maximizing utility, and markets for Arrow securities clear.

Proposition 2. *When Arrow securities are available, the unique equilibrium state price density is:*

$$\pi(\boldsymbol{x}) = C \cdot \exp\left(-\frac{W^*(\boldsymbol{x})}{\sum_{j=1}^{N} \alpha_j^{-1}}\right) \cdot f(\boldsymbol{x}) \tag{25}$$

where C is an arbitrary positive constant. Agents' asset demands are:

$$\theta_i(\boldsymbol{x}) = W_i^*(\boldsymbol{x}) - W_i^{Spot}(\boldsymbol{x}) \tag{26}$$

where:

$$W_i^*(\boldsymbol{x}) = \frac{\mathbb{E}\left[\exp\left(-\frac{W^*}{\sum_{j=1}^{N}\alpha_j^{-1}}\right) \cdot \left(W_i^{Spot} - \frac{\alpha_i^{-1}}{\sum_{j=1}^{N}\alpha_j^{-1}} W^*\right)\right]}{\mathbb{E}\left[\exp\left(-\frac{W^*}{\sum_{j=1}^{N}\alpha_j^{-1}}\right)\right]} + \frac{\alpha_i^{-1}}{\sum_{j=1}^{N}\alpha_j^{-1}} W^*(\boldsymbol{x}) \tag{27}$$

The equilibrium with Arrow securities is Pareto-efficient.

Technically, Proposition 2 is simply the classical welfare theorems, applied to our setting: if financial markets are complete, and the first-best outcome described in Proposition 1 is attainable through financial asset trading, then the first-best outcome must be an equilibrium.

Intuitively, spot markets and financial markets bring us from autarky to first-best efficiency in two stages, corresponding to the two efficiency conditions in Proposition 1. Spot market equilibrium implements the allocative efficiency condition (11), reallocating goods across consumers in each state of the world to optimally convert goods to wealth. This reduces the two-good problem to a one-good problem, where each consumer is effectively endowed with $W_i^{Spot}(\boldsymbol{x})$ dollars in state $\boldsymbol{x}$. However, spot markets do not efficiently distribute this wealth across agents.

Financial market equilibrium implements the risk-sharing condition (12), optimally sharing state-dependent shocks to aggregate wealth $W^*(\boldsymbol{x})$ across consumers. Note that equilibrium wealth with Arrow securities, (27), takes the form of first-best wealth (12), dividing aggregate wealth across agents according to their inverse risk aversions. Equilibrium Arrow security demands (26) are very simple: agents simply purchase the difference between their first-best wealths $W_i^*(\boldsymbol{x})$ and their spot-equilibrium wealths $W_i^{Spot}(\boldsymbol{x})$.

Financial markets are needed because spot market wealths $W_i^{Spot}(\boldsymbol{x})$ generally differ from first-best outcomes $W_i^*(\boldsymbol{x})$. In particular, risk aversions α_i influence first-best wealth allocations $W_i^*(\boldsymbol{x})$, but not spot market outcomes $W_i^{Spot}(\boldsymbol{x})$; clearly, spot market outcomes cannot implement first-best outcomes in general. Spot market outcomes occur after all inventory shock uncertainty $\boldsymbol{x}$ is realized, and thus clearly cannot allow consumers to share $\boldsymbol{x}$-related risk.[3]

5 Price Controls Can Be Pareto-Improving

When financial markets are missing, *price controls* in spot markets can be Pareto-improving, because the benefits they induce for risk-sharing can outweigh their costs for allocative efficiency.

We assume a policymaker can set a price ceiling p_{ceil}, which is constant and does not depend on the realization of $\boldsymbol{x}$. The price ceiling imposes an upper bound on market prices, which induces costless rationing when it is binding. Formally, suppose the unconstrained market clearing price exceeds p_{ceil}. Let $\mathcal{A}_B$ be the set of agents who purchase at p_{ceil}, that is, $q_i(p_{ceil}) > 0$, and let $\mathcal{A}_S$ be those agents with $q_i(p_{ceil}) < 0$. We assume all trade occurs at p_{ceil}, total trade volume equals total supply at p_{ceil}, and trade volume is rationed across buyers according to their relative demands at p_{ceil}. Formally,

$$q_i^{ceil}(p_{ceil}) = q_i(p_{ceil}) \ \ \forall i \in \mathcal{A}_S \tag{28}$$

$$q_i^{ceil}(p_{ceil}) = q_i(p_{ceil}) \left(-\frac{\sum_{i \in \mathcal{A}_S} q_i(p_{ceil})}{\sum_{i \in \mathcal{A}_B} q_i(p_{ceil})} \right) \ \ \forall i \in \mathcal{A}_B \tag{29}$$

Analogously, a price floor p_{floor} sets a lower bound on market prices; when binding, total trade volume equals total demand at p_{floor}, and quantity is rationed across sellers according to relative supply amounts:

$$q_i^{floor}(p_{floor}) = q_i(p_{floor}) \left(-\frac{\sum_{i \in \mathcal{A}_B} q_i(p_{floor})}{\sum_{i \in \mathcal{A}_S} q_i(p_{floor})} \right) \ \ \forall i \in \mathcal{A}_S \tag{30}$$

[3] Another way to see this is that, in a two-good spot market after $\boldsymbol{x}$ is realized, utility is ordinal rather than cardinal: a consumer's preferences are fully described by indifference curves between money and goods, which are traced out by (13). Expression (21) for $W_i^{Spot}(\boldsymbol{x})$ is thus valid regardless of what consumers' preferences over wealth are: (21) holds for any choice of CARA-utility risk aversions, or indeed any other classes of risk-averse preferences over wealth.

Evaluating welfare under price controls is straightforward. Given any inventory shocks $\boldsymbol{x}$, we solve for equilibrium prices, imposing price controls if they bind. We then use (28) and (29), and their analogs for price floors, to calculate equilibrium quantities; we then plug prices and quantities into (13) to calculate agents' equilibrium wealth levels, and thus CARA-utility levels, for any realization of $\boldsymbol{x}$. Expected utility under the price control regime is then calculated by integrating over the distribution of shocks.

Numerical Example. Suppose there are two consumers, with $\psi = 0$, $\alpha = 2$, $\kappa = 1$, and symmetrically distributed inventory shocks $x_1, x_2 \sim N(0, \sigma^2)$ with $\sigma^2 = 0.45$. Consumers are fully symmetric, so their ex-ante expected utilities are always identical. Figure 1 plots consumers' expected utility, in unconstrained spot markets (blue) and under varying levels of symmetric price controls (red), where we set a price ceiling $p_{ceil} = \bar{p}$ and a price floor $p_{floor} = -\bar{p}$. Price controls can be Pareto-improving: both consumers' expected utility is higher, for any value of $\bar{p}$ greater than around 1.2, relative to free spot markets.

The intuition for this result is illustrated in Figure 2. Spot markets alone fail to achieve perfect risk-sharing: consumers' marginal utilities are not equalized across states. Panel A plots the normalized difference in spot-equilibrium marginal utility, as a function of inventory shocks $\boldsymbol{x}$:

$$\Delta MU(\boldsymbol{x}) = \frac{MU_1(\boldsymbol{x}) - MU_2(\boldsymbol{x})}{MU_1(\boldsymbol{x}) + MU_2(\boldsymbol{x})} \tag{31}$$

In unconstrained spot market equilibrium, consumers with more extreme inventory shocks end up with lower wealth and higher marginal utility: 1's MU is greater towards the right and left, and 2's is greater upwards and downwards.[4] Thus, risk-sharing could improve, and aggregate welfare could increase, if wealth could be transferred from 2 to 1 on the right and left sides of the figure, and from 1 to 2 towards the top and bottom.

Panel B plots the net *wealth transfer* induced by the price control policy, defined as:

$$WealthTransfer(\boldsymbol{x}) = \frac{\left[W_1^{PC}(\boldsymbol{x}) - W_1^{Spot}(\boldsymbol{x})\right] - \left[W_2^{PC}(\boldsymbol{x}) - W_2^{Spot}(\boldsymbol{x})\right]}{2} \tag{32}$$

In words, (32) is a double-difference, measuring whether price controls increase 1's wealth more than they increase 2's wealth. The transfers induced by price controls are directionally consistent with improved risk-sharing. 1 is the net transfer recipient towards the left and right of the plot, and 2 is the net recipient towards the top and bottom. Intuitively, when x_1 is large and positive and x_2 is near 0, 1 is a net seller and 2 is a net buyer. Thus, a price

[4] Intuitively, spot market outcomes are qualitatively similar to no-trade outcomes in this case: if each consumer consumed their endowment, due to quadratic costs, consumers' marginal utilities of wealth would be decreasing in the magnitude of their inventory shocks.

floor tends to increase 1's welfare at the expense of 2, at the cost of some deadweight loss. This does not induce a net transfer ex-ante, because the reverse transfer occurs when x_2 is high and x_1 is near zero. Analogously, price ceilings transfer welfare towards 1 when x_1 is very negative and x_2 is near 0, and towards 2 in the reverse case.

Panel C of Figure 2 plots the *deadweight loss* from price controls, defined simply as the change in total social wealth:

$$DeadweightLoss\left(\boldsymbol{x}\right)=\frac{\left[W_1^{Spot}\left(\boldsymbol{x}\right)+W_2^{Spot}\left(\boldsymbol{x}\right)\right]-\left[W_1^{PC}\left(\boldsymbol{x}\right)+W_2^{PC}\left(\boldsymbol{x}\right)\right]}{2} \tag{33}$$

Deadweight loss is always positive, and tends to be greater when the net wealth transfer induced by price controls is larger. However, it is also in these regions that the risk-sharing benefits of price controls are greatest.

Figure 1: Price Controls and Expected Utility

This figure plots consumers' expected utility in unconstrained spot markets (blue) as well as under symmetric price floors and ceilings (red), where we set $p_{ceil}=-p_{floor}=\bar{p}\geq 0$. Higher values of $\bar{p}$, on the x-axis, thus correspond to looser price controls. For each $\bar{p}$ we calculate utility for a grid of $\boldsymbol{x}$-values, and numerically integrate to find expected utility.

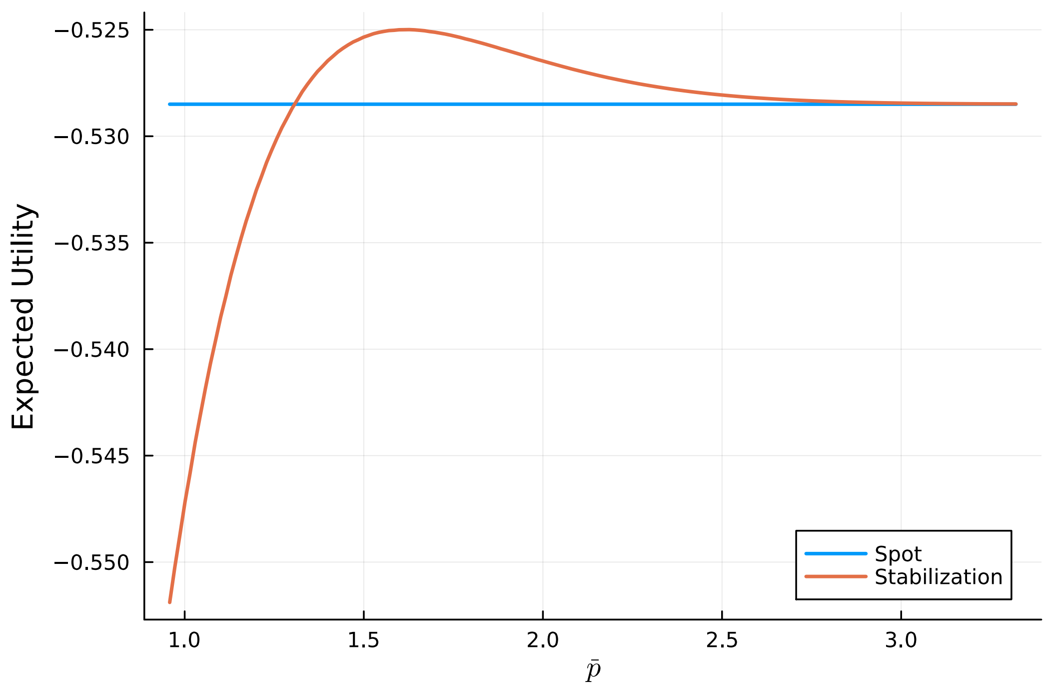

In models without uncertainty, price controls are *transfer* instruments: price ceilings transfer surplus to buyers, and price floors to sellers. Our model introduces a distinct *risk-sharing* role for price controls. In our stylized example, both agents are symmetric, and

Figure 2: Price Controls: Mechanisms

In each panel, consumers' inventory shocks x_1 and x_2 are shown on the x and y axes respectively. Panel A shows (31), the normalized difference between 1 and 2's marginal utility of wealth, induced by unconstrained spot market equilibrium outcomes. Panel B plots (32), the net wealth transfer from 2 to 1 induced by price controls, defined as the difference between 1's wealth gain under price controls relative to unconstrained spot market equilibrium, and 2's wealth gain. Price controls can improve risk sharing because they tend to transfer wealth in the direction of MU differences: wealth is transferred to 1 towards the right and left, where 1's marginal utility is higher, and to 2 towards the top and bottom, where 2's marginal utility is higher. Panel C plots (33), price control-induced deadweight loss, defined as total social wealth in spot market equilibrium minus total wealth under price controls. DWL is always positive, and is higher when price controls are more binding and induce larger transfers between agents. In both panels B and C, we consider a symmetric price control $p_{ceil} = -p_{floor} = 0.5$.

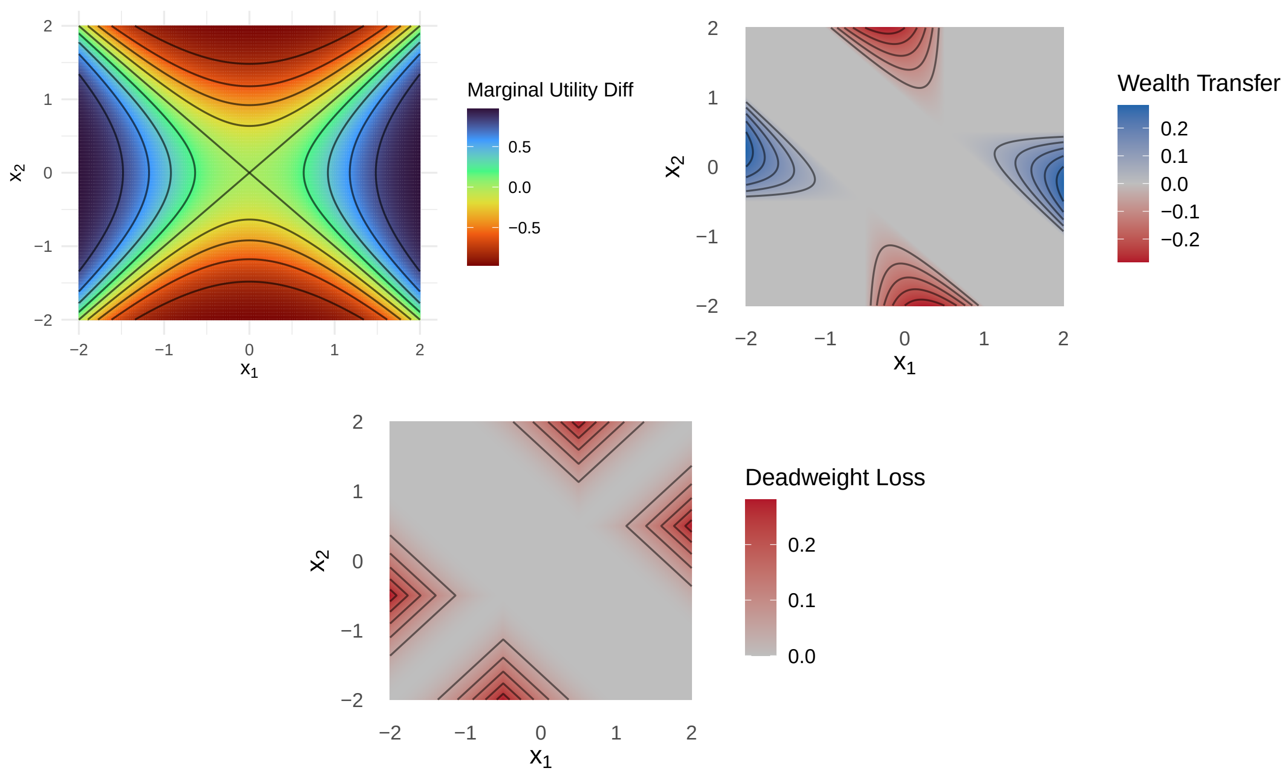

neither is a buyer or seller on average: price controls induce exactly offsetting transfers across uncertain states of the world, which can have the effect of improving both agents' ex-ante welfare. This is true even though price controls are harmful for allocative efficiency: social aggregate wealth unambiguously decreases whenever price controls are binding.

Market incompleteness is crucial for this result. Complete financial markets perfectly equalize agents' marginal utilities across states, leaving no further room for improvements in risk sharing. In incomplete financial markets, agents' marginal utilities may differ across states, leaving room for price stabilization policies to be welfare-improving.

Our results can be thought of as an instance of the "theorem of the second best": price controls are unambiguously welfare-reducing in the frictionless complete-market benchmark, but can be welfare-improving in the more realistic setting of incomplete financial markets.

6 Pointwise Optimal Price Controls

Next, we characterize optimal state-dependent price controls. While this is a much less realistic policy than state-independent price controls, the state-dependent problem is analytically simpler, and thus helps develop intuition about the comparative statics underlying the state-independent problem. Building on these results, Appendix A.5.2 then derives an analytical, though complex, first-order-condition for the optimal state-independent price controls; intuitively, the preferred fixed price controls optimally trades off positive and negative deviations from the pointwise optimal controls at each $\boldsymbol{x}$.

In the general case, we will allow state-dependent price floors $p_{floor}(\boldsymbol{x})$ and ceilings $p_{ceil}(\boldsymbol{x})$, though only one of the two will bind at any realization of $\boldsymbol{x}$. To present the main idea in a simplified manner, we first ignore price floors and consider the choice of a state-specific price ceiling policy, $p_{ceil}(\boldsymbol{x})$. Any Pareto-efficient price ceiling policy must maximize the sum of agents' weighted expected utilities:

$$\sum_{i=1}^{N} \lambda_i E\left[U_i\left(W_i\left(p_{ceil}(\boldsymbol{x}), \boldsymbol{x}\right)\right)\right] \tag{34}$$

where we write $W_i(p_{ceil}, \boldsymbol{x})$ to mean i's wealth when inventory shocks are $\boldsymbol{x}$ and the price ceiling is p_{ceil}. Since p_{ceil} is chosen separately for each $\boldsymbol{x}$, (34) simplifies to the pointwise optimization of p_{ceil} for each $\boldsymbol{x}$:

$$\max_{p_{ceil}} \sum_{i=1}^{N} \lambda_i U_i(W_i(p_{ceil}, \boldsymbol{x})) \tag{35}$$

where for simplicity we suppress dependence of p_{ceil} on $\boldsymbol{x}$. In the two-agent case we analyze in most of Section 5, (35) is simply:

$$\lambda_1 U_1\left(W_1\left(p_{ceil}, \boldsymbol{x}\right)\right)+\lambda_2 U_2\left(W_2\left(p_{ceil}, \boldsymbol{x}\right)\right) \tag{36}$$

for some weights $\lambda_1, \lambda_2 > 0$. Differentiating (36), if there is a binding optimal choice of p_{ceil}, it is characterized by the first-order condition:

$$0=\lambda_1 \alpha_1 e^{-\alpha_1 W_1\left(p_{ceil}, \boldsymbol{x}\right)} \cdot \frac{\partial W_1\left(p_{ceil}, \boldsymbol{x}\right)}{\partial p_{ceil}}+\lambda_2 \alpha_2 e^{-\alpha_2 W_2\left(p_{ceil}, \boldsymbol{x}\right)} \cdot \frac{\partial W_2\left(p_{ceil}, \boldsymbol{x}\right)}{\partial p} \tag{37}$$

Intuitively, (37) requires that small changes in p_{ceil} do not change weighted social welfare, (36), to first-order. Rearranging, we have:

$$\frac{-\frac{\partial W_1\left(p_{ceil}, \boldsymbol{x}\right)}{\partial p_{ceil}}}{\frac{\partial W_2\left(p_{ceil}, \boldsymbol{x}\right)}{\partial p_{ceil}}}=\frac{\lambda_2 \alpha_2 e^{-\alpha_2 W_2\left(p_{ceil}, \boldsymbol{x}\right)}}{\lambda_1 \alpha_1 e^{-\alpha_1 W_1\left(p_{ceil}, \boldsymbol{x}\right)}} \tag{38}$$

The LHS of (38) is the marginal inefficiency of p_{ceil} as a transfer tool: it measures how much 1's wealth decreases for each dollar that 2's wealth increases, as p_{ceil} increases. The RHS is the product of 2's relative utility weight $\lambda = \lambda_2/\lambda_1$, and the ratio of marginal utilities of wealth of 2 and 1. Intuitively, (38) states that, at the optimal p_{ceil}, the ratio between 1 and 2's marginal utilities of wealth (adjusted by λ) must be equal to the marginal DWL of p_{ceil}.

While (38) captures the intuition behind optimal statewise price control policies, the full problem is slightly more notationally complex because floors and ceilings will never be in use simultaneously. Thus, the optimal policy must be characterized casewise, depending on which of p_{ceil} and p_{floor} is binding. We state this in the following proposition, proved in Appendix A.5.

Proposition 3. *Consider the two-agent case with symmetric α, κ.*

For any state $\boldsymbol{x} = (x_1, x_2)$, define $M = \arg\max\{x_1, x_2\}$ and $m = \arg\min\{x_1, x_2\}$. Further define $h(\boldsymbol{x}) = W_M^{Spot}(\boldsymbol{x}) - W_m^{Spot}(\boldsymbol{x}) - (\ln \lambda_M - \ln \lambda_m)/\alpha$.

If $x_1 = x_2$ or $h(\boldsymbol{x}) = 0$, the spot market equilibrium price $p^{Spot}(\boldsymbol{x})$ characterizes the optimal price and neither price floor nor price ceiling is needed.

If $h(\boldsymbol{x}) > 0$, the optimal price ceiling policy $p_{ceil}(\boldsymbol{x})$ solves

$$\frac{\kappa\left(p-p^{Spot}(\boldsymbol{x})\right)+\frac{x_M-x_m}{2}}{3 \kappa\left(p-p^{Spot}(\boldsymbol{x})\right)+\frac{x_M-x_m}{2}}=\frac{x_M+\kappa(p-\psi)}{x_1+x_2+x_M+3 \kappa(p-\psi)}=\frac{\lambda_m}{\lambda_M} \cdot \frac{e^{-\alpha W_m(p, \boldsymbol{x})}}{e^{-\alpha W_M(p, \boldsymbol{x})}} \tag{39}$$

If $h(\boldsymbol{x}) < 0$, the optimal price floor policy $p_{floor}(x)$ solves

$$\frac{\kappa\left(p - p^{Spot}(\boldsymbol{x})\right) - \frac{x_M - x_m}{2}}{3\kappa\left(p - p^{Spot}(\boldsymbol{x})\right) - \frac{x_M - x_m}{2}} = \frac{x_m + \kappa(p - \psi)}{x_1 + x_2 + x_m + 3\kappa(p - \psi)} = \frac{\lambda_M}{\lambda_m} \cdot \frac{e^{-\alpha W_M(p,\boldsymbol{x})}}{e^{-\alpha W_m(p,\boldsymbol{x})}} \tag{40}$$

Proposition 3, intuitively, is simply the FOC in (38), substituting analytical expressions for the marginal transfer efficiency term on the LHS. Proposition 3 delivers two comparative statics results for the case $\lambda = 1$. The results are proved in Appendix A.5.1.

Firstly, an increase in risk aversion ($\alpha \uparrow$) leads to tighter optimal price controls: p_{ceil} decreases when it is binding, and p_{floor} increases when it is binding. Intuitively, risk aversion increases the marginal utility gaps implied by a given wealth gap, moving the ratio on the RHS of (38) away from 1; tighter price controls, inducing larger transfers and thus greater marginal inefficiency, are optimal as a result.

Secondly, an increase in ψ leads to a tighter optimal price ceiling and a looser optimal price floor: $p^{Spot} - p_{ceil}$ increases when p_{ceil} is binding, and $p_{floor} - p^{Spot}$ decreases when p_{floor} is binding. Intuitively, higher ψ delivers a good news for the consumer who has a big positive inventory shock (thus, the net supplier in the market) and decreases the marginal utility gap when the price is relatively low; while it delivers a bad news for the consumer who has a big negative shock and increases the marginal utility gap when the price is relatively high.

In the following claim, we derive state-dependent bounds on the optimal price controls.

Claim 1. Still consider the two-agent case with symmetric α, κ. We have

$$p_{ceil}(\boldsymbol{x}) \in \left(\psi - \frac{x_1 + x_2 + \max\{x_1, x_2\}}{3\kappa}, \psi - \frac{x_1 + x_2}{2\kappa}\right) \quad \text{and}$$
$$p_{floor}(\boldsymbol{x}) \in \left(\psi - \frac{x_1 + x_2}{2\kappa}, \psi - \frac{x_1 + x_2 + \min\{x_1, x_2\}}{3\kappa}\right) \tag{41}$$

where $p_{ceil}(\boldsymbol{x})$ indicates that the optimal policy is a binding price ceiling and $p_{floor}(\boldsymbol{x})$ indicates that it is a binding price floor.

The proof of Claim 1 is embedded in that of Proposition 3. The intuition behind Claim 1 is fairly simple: as price controls tighten, the deadweight loss they induce can become large to the point that further tightening price controls causes both consumers' utilities to decrease. As a result, for any given $\boldsymbol{x}$, the price controls which maximize a single consumer's utility is interior. The set of Pareto-efficient price control policies, which maximize the sum of consumers' weighted utilities, thus interpolate between these interior extremal points.

7 Conclusion

This paper has demonstrated that price controls can be Pareto-improving in the model of Han, Hu and Zhang (2026).

References

Borch, Karl. 1962. "Equilibrium in a Reinsurance Market." *Econometrica: Journal of the Econometric Society*, 30(3): 424–444.

Duffie, Darrell. 2010. *Dynamic asset pricing theory.* Princeton University Press.

Han, Baiyang, Junyi Hu, and Anthony Lee Zhang. 2026. "Markets for Price Risk." Working paper.

Wilson, Robert. 1968. "The theory of syndicates." *Econometrica: Journal of the Econometric Society*, 119–132.

Internet Appendix

A Omitted Proofs and Derivations

Appendices A.1, A.2, A.3, and A.4 are directly copied from Han, Hu and Zhang (2026). We include them here for expositional completeness.

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 2. For any constant A, define:

$$\tilde{\psi} \equiv \psi + A, \; \tilde{x}_i \equiv x_i + \kappa_i A \tag{42}$$

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

$$W_i = \tilde{\psi} \left(\tilde{x}_i + q_i \right) - \frac{\left(\tilde{x}_i + q_i \right)^2}{2\kappa_i} + C_i \left(A \right) \tag{43}$$

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

Claim 2 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 (42) is linear in A, this immediately implies that, for any set of original inventory shocks x_i which do not have 0 mean across consumers, we can find some A to normalize ψ and inventory shocks, which leads the resultant inventory shocks to have zero mean across consumers. This choice of A is simply:

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

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

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 (20) of Section 4, (5) is a natural normalization, because it implies that μ_i is equal to negative i's expected trade volume in spot markets.

A.1.1 Proof of Claim 2

Note that (42) 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 (43), 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 (21)

Copying (13), consumers' wealth is:

$$W_i = \psi\left(x_i + q_i\right) - \frac{\left(x_i + q_i\right)^2}{2\kappa_i} - pq_i \tag{44}$$

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 attain (21).

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:

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

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

The Lagrangian is:

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

The first-order condition is:

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

Implying simply that consumers' marginal rate of substitution between wealth and goods, $\frac{\partial W_i}{\partial y_i} = \psi - \frac{y_i}{\kappa_i}$, must be equated:

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

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.

A.4 Proof of Proposition 2

Spot market equilibrium endows agent i with $W_i^{Spot}(\boldsymbol{x})$ wealth in state $\boldsymbol{x}$. When agents can trade Arrow securities, markets are trivially complete, so the first welfare theorem implies that equilibrium allocations are Pareto-efficient. We use $W_i^*(\boldsymbol{x})$ to denote i's total equilibrium wealth in state $\boldsymbol{x}$: this is the sum of spot wealth $W_i^{Spot}(\boldsymbol{x})$ and any Arrow security payoffs $\theta_i(\boldsymbol{x})$. Agents' FOC for optimal security demand implies that the state price density is determined by agents' marginal utilities at $W_i^*(\boldsymbol{x})$:

$$\frac{\pi(\boldsymbol{x})}{\pi(\boldsymbol{x}')} = \frac{m(\boldsymbol{x}) \cdot f(\boldsymbol{x})}{m(\boldsymbol{x}') \cdot f(\boldsymbol{x}')} = \frac{U_i'(W_i^*(\boldsymbol{x})) \cdot f(\boldsymbol{x})}{U_i'(W_i^*(\boldsymbol{x}')) \cdot f(\boldsymbol{x}')}.$$

Using the representation of first-best wealth allocations in (12) of Proposition 1, we have:

$$\frac{U_i'(W_i^*(\boldsymbol{x})) \cdot f(\boldsymbol{x})}{U_i'(W_i^*(\boldsymbol{x}')) \cdot f(\boldsymbol{x}')} = \frac{\exp\left(-\frac{W^*(\boldsymbol{x})}{\sum_{j=1}^N \alpha_j^{-1}}\right) \cdot f(\boldsymbol{x})}{\exp\left(-\frac{W^*(\boldsymbol{x}')}{\sum_{j=1}^N \alpha_j^{-1}}\right) \cdot f(\boldsymbol{x}')} \tag{45}$$

which gives (25). Notice that, under CARA utility, all Pareto efficient allocations imply the same state-price density: the constant terms in (12) fall out of the ratio in (45).

To calculate equilibrium Arrow security demands, note that spot market equilibrium endows i with wealth $W_i^{Spot}(\boldsymbol{x})$ in state $\boldsymbol{x}$, and Pareto-efficient wealth allocations $W_i^*(\boldsymbol{x})$ have the form in (12) of Proposition 1. In order for $\theta_i(\boldsymbol{x})$ to induce Pareto-efficient wealth allocations, we must have, for each i:

$$\theta_i(\boldsymbol{x}) = W_i^*(\boldsymbol{x}) - W_i^{Spot}(\boldsymbol{x}) = C_i + \frac{\alpha_i^{-1}}{\sum_{j=1}^N \alpha_j^{-1}} W^*(\boldsymbol{x}) - W_i^{Spot}(\boldsymbol{x}). \tag{46}$$

for some C_i. We can find C_i using the budget constraint (23), substituting (25) and (46):

$$C\int\left(C_i+\frac{\alpha_i^{-1}}{\sum_{j=1}^N\alpha_j^{-1}}W^*(\boldsymbol{x})-W_i^{Spot}(\boldsymbol{x})\right)\cdot\exp\left(-\frac{W^*(\boldsymbol{x})}{\sum_{j=1}^N\alpha_j^{-1}}\right)f(\boldsymbol{x})\,d\boldsymbol{x}=0$$

Solving, we have:

$$C_i=\frac{\mathbb{E}\left[\exp\left(-\frac{W^*}{\sum_{j=1}^N\alpha_j^{-1}}\right)\cdot\left(W_i^{Spot}-\frac{\alpha_i^{-1}}{\sum_{j=1}^N\alpha_j^{-1}}W^*\right)\right]}{\mathbb{E}\left[\exp\left(-\frac{W^*}{\sum_{j=1}^N\alpha_j^{-1}}\right)\right]}$$

This gives (27).

A.5 Proof of Proposition 3

We first derive the wealth levels for any given price p for any state $\boldsymbol{x}=(x_1,x_2)$ under the rationing policy given in Section 5.

For simplicity we suppress dependence of demand on $\boldsymbol{x}$. The demand for consumer i is given by $q_i(p)=-\kappa(p-\psi)-x_i$. Market clears at $p^{Spot}(\boldsymbol{x})=\psi-\frac{x_1+x_2}{2\kappa}$. If $x_1=x_2$, then we have $q_1(p)=q_2(p)$ for every p, which implies trade never occurs in the good market. This is an edge case where no price control policy is needed. Below we focus on $x_1\neq x_2$. We have four cases.

- If $p\leq\psi-\frac{x_M}{\kappa}$, $q_m(p)>q_M(p)\geq 0$. No trade occurs and we have $W_i(p,\boldsymbol{x})=W_i^0=\psi x_i-\frac{x_i^2}{2\kappa}$, the wealth level in autarky.
- If $\psi-\frac{x_M}{\kappa}<p\leq\psi-\frac{x_1+x_2}{2\kappa}$, $q_M(p)<0<-q_M(p)\leq q_m(p)$ and p is a binding price ceiling. Consumer M's demand is unchanged while that of consumer m needs to be rationed: $q_M^{ceil}(p_{ceil})=-q_m^{ceil}(p_{ceil})=-\kappa(p_{ceil}-\psi)-x_M$. The wealth is given by

$$W_M(p_{ceil},\boldsymbol{x})=p_{ceil}x_M+\frac{\kappa(\psi-p_{ceil})^2}{2}\quad\text{and}$$
$$W_m(p_{ceil},\boldsymbol{x})=p_{ceil}x_m-\frac{3\kappa(\psi-p_{ceil})^2}{2}+2(\psi-p_{ceil})(x_1+x_2)-\frac{(x_1+x_2)^2}{2\kappa}$$

- If $\psi-\frac{x_1+x_2}{2\kappa}<p<\psi-\frac{x_m}{\kappa}$, $q_M(p)<0<q_m(p)<-q_M(p)$ and p is a binding price floor. Consumer m's demand is unchanged while that of consumer M needs to be rationed:

$-q_M^{floor}(p_{floor}) = q_m^{floor}(p_{floor}) = -\kappa(p_{floor} - \psi) - x_m$. The wealth is given by

$$W_M(p_{floor}, \boldsymbol{x}) = p_{floor}x_M - \frac{3\kappa(\psi - p_{floor})^2}{2} + 2(\psi - p_{floor})(x_1 + x_2) - \frac{(x_1 + x_2)^2}{2\kappa} \quad \text{and}$$

$$W_m(p_{floor}, \boldsymbol{x}) = p_{floor}x_m + \frac{\kappa(\psi - p_{floor})^2}{2}$$

- If $p \geq \psi - \frac{x_m}{\kappa}$, $q_M(p) < q_m(p) \leq 0$. No trade occurs and we have $W_i(p, \boldsymbol{x}) = W_i^0 = \psi x_i - \frac{x_i^2}{2\kappa}$.

Note that $W_i(p, \boldsymbol{x})$ peaks at $\psi - \frac{x_1+x_2+x_{-i}}{3\kappa}$ and is decreasing on both sides. Hence, we get the bound in (41) and conclude Claim 1.

Now that we get the wealth for any p and $\boldsymbol{x}$, we can plug it in the FOC in (37).

If the optimal price control policy lies within $\psi - \frac{x_1+x_2+x_M}{3\kappa}$ and $\psi - \frac{x_1+x_2}{2\kappa}$, it is a binding price ceiling. The optimal price ceiling policy $p_{ceil}(\boldsymbol{x})$ solves

$$0 = \lambda_M(x_M + \kappa(p - \psi))e^{-\alpha W_M(p,\boldsymbol{x})} - \lambda_m(x_1 + x_2 + x_M + 3\kappa(p - \psi))e^{-\alpha W_m(p,\boldsymbol{x})}. \tag{47}$$

Rearranging, we have

$$\frac{x_M + \kappa(p - \psi)}{x_1 + x_2 + x_M + 3\kappa(p - \psi)} = \frac{\lambda_m}{\lambda_M} \cdot \frac{e^{-\alpha W_m(p,\boldsymbol{x})}}{e^{-\alpha W_M(p,\boldsymbol{x})}} = e^{\alpha(W_M(p,\boldsymbol{x}) - W_m(p,\boldsymbol{x})) - (\ln\lambda_M - \ln\lambda_m)} \tag{48}$$

Within $\left(\psi - \frac{x_1+x_2+x_M}{3\kappa}, \psi - \frac{x_1+x_2}{2\kappa}\right)$, the LHS of (48) is decreasing in p and achieves 1 at $p = \psi - \frac{x_1+x_2}{2\kappa}$, while the RHS of (48) is increasing in p. Hence, if $h(\boldsymbol{x}) = W_M^{Spot}(\boldsymbol{x}) - W_m^{Spot}(\boldsymbol{x}) - (\ln\lambda_M - \ln\lambda_m)/\alpha \leq 0$, the RHS is always smaller than the LHS. This implies that a binding price ceiling policy is not optimal and we should look for a price floor policy instead. If $h(\boldsymbol{x}) > 0$, FOC in (48) has a unique solution within $\left(\psi - \frac{x_1+x_2+x_M}{3\kappa}, \psi - \frac{x_1+x_2}{2\kappa}\right)$,which gives the optimal price ceiling policy $p_{ceil}(\boldsymbol{x})$.

If the optimal price control policy lies within $\psi - \frac{x_1+x_2}{2\kappa}$ and $\psi - \frac{x_1+x_2+x_m}{3\kappa}$, it is a binding price floor. The optimal price floor policy $p_{floor}(\boldsymbol{x})$ solves

$$0 = -\lambda_M(x_1 + x_2 + x_m + 3\kappa(p - \psi))e^{-\alpha W_M(p,\boldsymbol{x})} + \lambda_m(x_m + \kappa(p - \psi))e^{-\alpha W_m(p,\boldsymbol{x})} \tag{49}$$

Rearranging, we have

$$\frac{x_m + \kappa (p - \psi)}{x_1 + x_2 + x_m + 3\kappa (p - \psi)} = \frac{\lambda_M}{\lambda_m} \cdot \frac{e^{-\alpha W_M(p,\boldsymbol{x})}}{e^{-\alpha W_m(p,\boldsymbol{x})}} = e^{-\alpha(W_M(p,\boldsymbol{x}) - W_m(p,\boldsymbol{x})) - (\ln \lambda_M - \ln \lambda_m)} \tag{50}$$

Within $\left(\psi - \frac{x_1+x_2}{2\kappa}, \psi - \frac{x_1+x_2+x_m}{3\kappa}\right)$, the LHS of (50) is increasing in p and achieves 1 at $p = \psi - \frac{x_1+x_2}{2\kappa}$, while the RHS of (50) is decreasing in p. If $h(\boldsymbol{x}) \geq 0$, the RHS is always smaller than the LHS. This implies that a binding price floor policy is not optimal and we should look for a price ceiling policy instead. If $h(\boldsymbol{x}) < 0$, FOC in (50) has a unique solution within $\left(\psi - \frac{x_1+x_2}{2\kappa}, \psi - \frac{x_1+x_2+x_m}{3\kappa}\right)$,which gives the optimal price floor policy $p_{floor}(\boldsymbol{x})$.

In all, if $h(\boldsymbol{x}) = 0$, the spot market equilibrium price $p^{Spot}(\boldsymbol{x})$ solves the FOC in (37) and characterizes the optimal price. This is also an edge case where no price control policy is needed, like $x_1 = x_2$. If $h(\boldsymbol{x}) > 0$, the optimal policy is a binding price ceiling which solves (48). If $h(\boldsymbol{x}) < 0$, the optimal policy is a binding price floor which solves (50). This concludes the proof of Proposition 3.

A.5.1 Proof of Comparative Statics of Pointwise Optimal Price Controls

In this section, we derive the two comparative statics results with $\lambda_1 = \lambda_2$.

Risk aversion parameter. We begin with the risk aversion parameter α. When $\lambda_1 = \lambda_2$, the sign of $h(\boldsymbol{x})$ is irrelevant of α. That is, if state $\boldsymbol{x}$ associates with an optimal price ceiling (floor) policy at some α, then it must associate with an optimal price ceiling (floor) policy at any level of risk aversion. If $h(\boldsymbol{x}) > 0$, taking derivatives w.r.t. α for both sides of (48), we get

$$\frac{\partial LHS}{\partial p} \cdot \frac{\partial p}{\partial \alpha} = \frac{\partial RHS}{\partial p} \cdot \frac{\partial p}{\partial \alpha} + \frac{\partial RHS}{\partial \alpha}.$$

Rearrange, we get

$$\left(\frac{\partial LHS}{\partial p} - \frac{\partial RHS}{\partial p}\right) \cdot \frac{\partial p}{\partial \alpha} = \frac{\partial RHS}{\partial \alpha}.$$

We have that the LHS is decreasing in p and the RHS is increasing in p. Also, $\partial RHS/\partial \alpha > 0$ at $p_{ceil}(\boldsymbol{x})$. Hence, $\partial p/\partial \alpha < 0$. Note that we have a binding price ceiling for this case and the negative derivatives implies a tighter optimal price control when α increases.

If $h(\boldsymbol{x}) < 0$, we take derivatives w.r.t. α for both sides of (50). Following similar analysis as above, we get $\partial p/\partial \alpha > 0$. Note that we have a binding price floor for this case and the positive derivatives implies a tighter optimal price control when α increases. In all, we concludes that an increase in risk aversion ($\alpha \uparrow$) leads to tighter optimal price controls: p_{ceil} decreases when it is binding, and p_{floor} increases when it is binding.

Value parameter ψ. We now turn to the parameter ψ. Note that the sign of $h(\boldsymbol{x})$ changes with ψ. Here, we focus on the case where changes in ψ does not flip the sign of $h(\boldsymbol{x})$: that is, if state $\boldsymbol{x}$ associates with an optimal price ceiling (floor) policy at some ψ, then it also associates with an optimal price ceiling (floor) policy when we consider changes in ψ. Also, when ψ changes, the spot market equilibrium price $p^{Spot}(\boldsymbol{x})$ changes simultaneously, so we consider the gap between the optimal price policy and $p^{Spot}(\boldsymbol{x})$ instead: $p^{gap}(\boldsymbol{x}) \equiv p(\boldsymbol{x}) - p^{Spot}(\boldsymbol{x})$. If $h(\boldsymbol{x}) > 0$, we can rewrite (48) as

$$\frac{\kappa p^{gap} + \frac{x_M - x_m}{2}}{3\kappa p^{gap} + \frac{x_M - x_m}{2}} = e^{\alpha(W_M(p^{gap}, \boldsymbol{x}) - W_m(p^{gap}, \boldsymbol{x}))} \tag{51}$$

where

$$W_M(p^{gap}, \boldsymbol{x}) = p^{gap} x_M + \left(\psi - \frac{x_1 + x_2}{2\kappa}\right) x_M + \frac{\kappa}{2}\left(p^{gap} - \frac{x_1 + x_2}{2\kappa}\right)^2 \quad \text{and}$$

$$W_m(p^{gap}, \boldsymbol{x}) = p^{gap} x_m + \left(\psi - \frac{x_1 + x_2}{2\kappa}\right) x_m - \frac{3\kappa}{2}\left(p^{gap} - \frac{x_1 + x_2}{2\kappa}\right)^2 - 2p^{gap}(x_1 + x_2) + \frac{(x_1 + x_2)^2}{2\kappa}.$$

Taking derivatives w.r.t. ψ for both sides of (51), we get

$$\frac{\partial LHS}{\partial p^{gap}} \cdot \frac{\partial p^{gap}}{\partial \psi} + \frac{\partial LHS}{\partial \psi} = \frac{\partial RHS}{\partial p^{gap}} \cdot \frac{\partial p^{gap}}{\partial \psi} + \frac{\partial RHS}{\partial \psi}.$$

Rearranging, we get

$$\left(\frac{\partial LHS}{\partial p^{gap}} - \frac{\partial RHS}{\partial p^{gap}}\right) \cdot \frac{\partial p^{gap}}{\partial \psi} = \frac{\partial RHS}{\partial \psi}.$$

We have that the LHS is decreasing in p (hence also p^{gap}) and the RHS is increasing in p (hence also p^{gap}). Also, $\partial RHS/\partial \psi = \alpha(x_M - x_m) e^{\alpha(W_M(p^{gap}, \boldsymbol{x}) - W_m(p^{gap}, \boldsymbol{x}))} > 0$. Hence, $\partial p^{gap}/\partial \psi < 0$. That is we have a tighter optimal price ceiling when ψ increases.

If $h(\boldsymbol{x}) < 0$, following similar analysis as above, we get $\partial p^{gap}/\partial \psi < 0$. That is, we have a looser optimal price floor.

Inventory Shock x_M and x_m. We then turn to the inventory shock x_M and x_m. Each time we change one of them and hold the other one fixed.

We first consider decreasing x_m while holding x_M fixed. Note that the sign of $h(\boldsymbol{x})$ changes with x_m. Here, we focus on the case where x_m changes within the range such that $h(\boldsymbol{x}) > 0$. That is an binding price ceiling policy is optimal. This corresponds to Figure 2 towards the left and bottom. Again, when x_m changes, the spot market equilibrium price $p^{Spot}(\boldsymbol{x})$ changes simultaneously, so we consider the gap $p^{gap}(\boldsymbol{x})$. Taking derivatives w.r.t. x_m for both sides of (51), we get

$$\frac{\partial LHS}{\partial p^{gap}} \cdot \frac{\partial p^{gap}}{\partial x_m} + \frac{\partial LHS}{\partial x_m} = \frac{\partial RHS}{\partial p^{gap}} \cdot \frac{\partial p^{gap}}{\partial x_m} + \frac{\partial RHS}{\partial x_m}.$$

Rearrange, we get

$$\left(\frac{\partial LHS}{\partial p^{gap}} - \frac{\partial RHS}{\partial p^{gap}}\right) \cdot \frac{\partial p^{gap}}{\partial x_m} = \frac{\partial RHS}{\partial x_m} - \frac{\partial LHS}{\partial x_m}.$$

We have that the LHS is decreasing in p^{gap} and the RHS is increasing in p^{gap}. Also, $\partial LHS/\partial x_m > 0$ given $p^{gap} < 0$, and

$$\begin{aligned}\frac{\partial RHS}{\partial x_m} = \alpha\left(-p^{gap} - \psi + \frac{x_m}{k}\right) e^{\alpha(W_M(p^{gap},\boldsymbol{x}) - W_m(p^{gap},\boldsymbol{x}))} = \\ \alpha\left(-\left(p^{gap} + \frac{x_M - x_m}{2\kappa}\right) - p^{Spot}(\boldsymbol{x})\right) e^{\alpha(W_M(p^{gap},\boldsymbol{x}) - W_m(p^{gap},\boldsymbol{x}))} < \\ \alpha\left(-\frac{x_M - x_m}{3\kappa} - p^{Spot}(\boldsymbol{x})\right) e^{\alpha(W_M(p^{gap},\boldsymbol{x}) - W_m(p^{gap},\boldsymbol{x}))} < 0\end{aligned}$$

given that $h(\boldsymbol{x}) > 0$ is equivalent to $p^{Spot}(\boldsymbol{x}) > 0$ under $\lambda_1 = \lambda_2$. Hence, $\partial p^{gap}/\partial x_m > 0$. That is we have a tighter optimal price ceiling policy when x_m decreases holding x_M fixed.

We then consider increasing x_M while holding x_m fixed. Again, the sign of $h(\boldsymbol{x})$ changes with x_M. This time, we focus on the case where x_M changes within the range such that $h(\boldsymbol{x}) < 0$. That is an binding price floor policy is optimal. This corresponds to Figure 2 towards the right and top. If $h(\boldsymbol{x}) < 0$, we can rewrite (50) as

$$\frac{\kappa p^{gap} - \frac{x_M - x_m}{2}}{3\kappa p^{gap} - \frac{x_M - x_m}{2}} = e^{\alpha(W_m(p^{gap},\boldsymbol{x}) - W_M(p^{gap},\boldsymbol{x}))} \tag{52}$$

where

$$W_M\left(p^{gap},\boldsymbol{x}\right)=p^{gap}x_M+\left(\psi-\frac{x_1+x_2}{2\kappa}\right)x_M-\frac{3\kappa}{2}\left(p^{gap}-\frac{x_1+x_2}{2\kappa}\right)^2-2p^{gap}\left(x_1+x_2\right)+\frac{\left(x_1+x_2\right)^2}{2\kappa}\quad\text{and}$$
$$W_m\left(p^{gap},\boldsymbol{x}\right)=p^{gap}x_m+\left(\psi-\frac{x_1+x_2}{2\kappa}\right)x_m+\frac{\kappa}{2}\left(p^{gap}-\frac{x_1+x_2}{2\kappa}\right)^2.$$

Taking derivatives w.r.t. x_M for both sides of (52), we get

$$\frac{\partial LHS}{\partial p^{gap}}\cdot\frac{\partial p^{gap}}{\partial x_M}+\frac{\partial LHS}{\partial x_M}=\frac{\partial RHS}{\partial p^{gap}}\cdot\frac{\partial p^{gap}}{\partial x_M}+\frac{\partial RHS}{\partial x_M}.$$

Rearrange, we get

$$\left(\frac{\partial LHS}{\partial p^{gap}}-\frac{\partial RHS}{\partial p^{gap}}\right)\cdot\frac{\partial p^{gap}}{\partial x_M}=\frac{\partial RHS}{\partial x_M}-\frac{\partial LHS}{\partial x_M}.$$

We have that the LHS is increasing in p (hence also p^{gap}) and the RHS is decreasing in p (hence also p^{gap}). Also, $\partial LHS/\partial x_M<0$ given $p^{gap}>0$, and

$$\begin{aligned}\frac{\partial RHS}{\partial x_M}&=\alpha\left(-p^{gap}-\psi+\frac{x_M}{k}\right)e^{\alpha\left(W_m\left(p^{gap},\boldsymbol{x}\right)-W_M\left(p^{gap},\boldsymbol{x}\right)\right)}=\\&\alpha\left(-\left(p^{gap}-\frac{x_M-x_m}{2\kappa}\right)-p^{Spot}\left(\boldsymbol{x}\right)\right)e^{\alpha\left(W_m\left(p^{gap},\boldsymbol{x}\right)-W_M\left(p^{gap},\boldsymbol{x}\right)\right)}>\\&\alpha\left(\frac{x_M-x_m}{3\kappa}-p^{Spot}\left(\boldsymbol{x}\right)\right)e^{\alpha\left(W_m\left(p^{gap},\boldsymbol{x}\right)-W_M\left(p^{gap},\boldsymbol{x}\right)\right)}>0\end{aligned}$$

given that $h\left(\boldsymbol{x}\right)<0$ is equivalent to $p^{Spot}\left(\boldsymbol{x}\right)<0$ under $\lambda_1=\lambda_2$. Hence, $\partial p^{gap}/\partial x_M>0$. That is we have a tighter optimal price floor policy when x_M increases holding x_m fixed.

A.5.2 FOC for Fixed Optimal Price Controls

We analyze price controls in a two-agent economy with symmetric parameters: identical α, κ, and independent, mean-zero inventory shocks $x_1,x_2\sim N\left(0,\sigma^2\right)$. We consider the price floor and ceiling policy: (p_{ceil},p_{floor}). Define the welfare gain from this policy as a function of the pair (p_{ceil},p_{floor}):

$$\Delta\left(p_{ceil},p_{floor}\right)=E\left[U_1\left(W_1^{floor=p_{floor},ceiling=p_{ceil}}\left(\boldsymbol{x}\right)\right)\right]-E\left[U_1\left(W_1^{Spot}\left(\boldsymbol{x}\right)\right)\right].$$

Then we can find the optimal (p_{ceil},p_{floor}) by looking at the first-order conditions of Δ.

The explicit form of $\Delta\left(p_{ceil},p_{floor}\right)$ is given by:

$$
\begin{aligned}
\Delta\left(\bar{p}\right)=&-\iint_{x_1\in A,\,x_2\in A}\exp\left(-\alpha\left(\psi x_1-\frac{x_1^2}{2\kappa}\right)\right)dF\left(x_1\right)dF(x_2)\\
&-\iint_{x_1\in A^c,\,x_1+x_2\in B}\exp\left(-\alpha\left(\frac{1}{2}\kappa\tilde{p}_{floor}^2+p_{floor}x_1\right)\right)dF\left(x_1\right)dF(x_2)\\
&-\iint_{x_2\in A^c,\,x_1+x_2\in B}\exp\left(\alpha\left(\frac{3}{2}\kappa\tilde{p}_{floor}^2+2\tilde{p}_{floor}\left(x_1+x_2\right)-p_{floor}x_1+\frac{\left(x_1+x_2\right)^2}{2\kappa}\right)\right)dF\left(x_1\right)dF(x_2)\\
&+\iint_{x_1+x_2\in B}\exp\left(-\alpha\left(\psi x_1+\frac{1}{8}\frac{\left(x_1+x_2\right)^2}{\kappa}-x_1\frac{\left(x_1+x_2\right)}{2\kappa}\right)\right)dF\left(x_1\right)dF(x_2)\\
&-\iint_{x_1\in C,\,x_2\in C}\exp\left(-\alpha\left(\psi x_1-\frac{x_1^2}{2\kappa}\right)\right)dF\left(x_1\right)dF(x_2)\\
&-\iint_{x_1\in C^c,\,x_1+x_2\in D}\exp\left(-\alpha\left(\frac{1}{2}\kappa\tilde{p}_{ceil}^2+p_{ceil}x_1\right)\right)dF\left(x_1\right)dF(x_2)\\
&-\iint_{x_2\in C^c,\,x_1+x_2\in D}\exp\left(\alpha\left(\frac{3}{2}\kappa\tilde{p}_{ceil}^2+2\tilde{p}_{ceil}\left(x_1+x_2\right)-p_{ceil}x_1+\frac{\left(x_1+x_2\right)^2}{2\kappa}\right)\right)dF\left(x_1\right)dF(x_2)\\
&+\iint_{x_1+x_2\in D}\exp\left(-\alpha\left(\psi x_1+\frac{1}{8}\frac{\left(x_1+x_2\right)^2}{\kappa}-x_1\frac{\left(x_1+x_2\right)}{2\kappa}\right)\right)dF\left(x_1\right)dF(x_2),
\end{aligned}
$$

where $\tilde{p}_{floor}=p_{floor}-\psi$, $\tilde{p}_{ceil}=p_{ceil}-\psi$ and $A=\left\{x>-\kappa\tilde{p}_{floor}\right\}$, $B=\left\{x>-2\kappa\tilde{p}_{floor}\right\}$, $C=\left\{x<-\kappa\tilde{p}_{ceil}\right\}$, $D=\left\{x<-2\kappa\tilde{p}_{ceil}\right\}$.

Taking derivatives w.r.t. p_{floor} we get

$$
\frac{\partial\Delta}{\partial p_{floor}}=\alpha\iint_{x_1<-\kappa\left(p_{floor}-\psi\right),\,x_1+x_2>-2\kappa\left(p_{floor}-\psi\right)}H\left(x_1,x_{2,}p_{floor}\right)dF\left(x_1\right)dF(x_2)\quad,
$$

where

$$
\begin{aligned}
H\left(x_1,x_{2,}p\right)=&\left(\kappa\left(p-\psi\right)+x_1\right)\exp\left(-\alpha\left(\frac{1}{2}\kappa\left(p-\psi\right)^2+px_1\right)\right)\\
&-\left(3\kappa\left(p-\psi\right)+\left(2x_1+x_2\right)\right)\exp\left(\alpha\left(\frac{3}{2}\kappa\left(p-\psi\right)^2+2\left(p-\psi\right)\left(x_1+x_2\right)-px_2+\frac{\left(x_1+x_2\right)^2}{2\kappa}\right)\right)
\end{aligned}
$$

We have $x_1<x_2$ in the entire integral region. Thus, $H(x_1,x_2,p_{floor})$ is the same as the RHS of the first-order condition (49) for point-wise optimal price floor policy under $\lambda_1=\lambda_2=1$. We can interpret $\partial\Delta/\partial p_{floor}$ as the integral of the point-wise first-order condition within the entire binding region.

Taking derivatives w.r.t. p_{ceil} we get

$$\frac{\partial \Delta}{\partial p_{ceil}} = \alpha \iint_{x_1 > -\kappa(p_{ceil}-\psi),\, x_1+x_2 < -2\kappa(p_{ceil}-\psi)} H\left(x_1, x_{2,} p_{floor}\right) dF\left(x_1\right) dF(x_2) \quad .$$

This time, we have $x_1 > x_2$ in the entire integral region. Thus, $H(x_1, x_2, p_{ceil})$ is the same as the RHS of the first-order condition (47) for point-wise optimal price ceiling policy under $\lambda_1 = \lambda_2 = 1$. We can interpret $\partial\Delta/\partial p_{ceil}$ as the integral of the point-wise first-order condition within the entire binding region.

An interior optimal price floor policy satisfy the first-order condition $\partial\Delta/\partial p_{floor} = 0$, while an interior optimal price ceiling policy satisfy the second-order condition $\partial\Delta/\partial p_{ceil} = 0$. In addition, if we consider symmetric price floor and ceiling, where we set $p_{ceil} - \psi = \psi - p_{floor} = \bar{p} \geq 0$. Then an interior optimal symmetric price control policy satisfy the first-order condition:

$$\begin{aligned} 0 = \frac{d\Delta}{d\bar{p}} = \frac{\partial \Delta}{\partial p_{floor}} \cdot \frac{dp_{floor}}{d\bar{p}} + \frac{\partial \Delta}{\partial p_{ceil}} \cdot \frac{dp_{ceil}}{d\bar{p}} = -\frac{\partial \Delta\left(\psi+\bar{p}, \psi-\bar{p}\right)}{\partial p_{floor}} + \frac{\partial \Delta\left(\psi+\bar{p}, \psi-\bar{p}\right)}{\partial p_{ceil}} \\ = -\alpha \iint_{x_1 < \kappa\bar{p},\, x_1+x_2 > 2\kappa\bar{p}} H\left(x_1, x_{2,} \psi-\bar{p}\right) dF\left(x_1\right) dF(x_2) + \\ \alpha \iint_{x_1 > -\kappa\bar{p},\, x_1+x_2 < -2\kappa\bar{p}} H\left(x_1, x_{2,} \psi+\bar{p}\right) dF\left(x_1\right) dF(x_2). \end{aligned}$$

Though we could get clean comparative statics for the point-wise price policy case, when p_{floor} and p_{ceil} enter the limits of the integration, things get complex. For example, an increase in α asks for a tighter point-wise optimal policy control, but in the fixed price control setting, a tighter policy (a lower p_{ceil} or a higher p_{floor}) also enlarges the regions where the price ceiling or floor is binding. Hence, we see the U shape in the comparative statics figure for the optimal fixed price control policy.