Shortages and Runs

Dean Jens

January 12, 2018

Abstract

In classical economic models, prices are expected to adjust rapidly to clear markets, but in real-world markets, prices are frequently "sticky", especially in the short run; excess demand may result in shortages. If agents anticipate this, a dynamic market in which deep preferences and supply are stationary may admit equilibria in which shortages take place, even if prices are such that the market would also clear without any rationing. A model is constructed in which it is shown that rationing is more likely to obtain for goods that are storable and "necessity", rather than "luxury", items. These are also the kinds of goods for which Cavallo, Cavallo, and Rigobon (2014) observe extended shortages in the event of supply disruptions after two large earthquakes. In situations without supply disruptions, the rationing equilibria are Pareto inferior to the classical equilibrium. I show that if agents believe that price will rise in response to shortages, the Pareto dominated equilibria go away, as long as the price is sufficiently responsive to the shortage.[1]

[1] The latest version of this paper is at http://eden.rutgers.edu/~djens/jobmarket/JensJMP.pdf.

1 Introduction

In 1973, a congressman from a forested district attempted to raise concerns that the government had made insufficient efforts to procure paper products. On December 19, 1973, Johnny Carson's Tonight Show picked up on one of his releases, and made some jokes about an impending toilet paper shortage.[2] In response, shoppers the next day bought out the entire supply of toilet paper at many stores; the empty store shelves confirmed to other shoppers that, indeed, there was a shortage of toilet paper, and as soon as the store resupplied the run continued.

In times of actual or perceived shortage, it is common to want to stockpile something just in case, even if it isn't needed right now. This, naturally, exacerbates the shortage, and may even be self-fulfilling. Gasoline is perhaps a more common example than toilet paper; if nobody tried to fill their gas tank precautionarily, there would be no shortage, but if I need gas tomorrow, and I'm afraid people who don't need it tomorrow will take what's available then I'm helping create the shortage today. In the event of a temporary supply disruption, this effect can prolong the shortage after supply comes back on line; if agents could coordinate, at that point everyone would be able to buy and consume at (or near) the "normal" level, but an expectation that the shortage will continue while consumers replenish their depleted stocks brings about the behavior that prolongs it.

This is related to the idea of liquidity premium; it some way it is its flip side. If there is a liquid market in which to sell an item, it can gain a liquidity premium: if I believe there is some chance that I will have a sudden need for any random other thing for which I can trade it (possibly indirectly), then

[2] The story is told at the priceonomics website, among other places; there is a youtube video of some of Carson's remarks.

by buying the item I'm buying that service as well as the item's use value itself. On the other hand, if there's a liquid market in which to *buy* an item, that can *lower* the value of the asset — as is perhaps more starkly noted when that market *doesn't* exist. The value of having some (storable) item now that may be hard, expensive, or even impossible to buy when I might suddenly need it includes a premium for the chance that I will need it a great deal at some point in the future and be unable to acquire it easily then. If I can get it when I need it, there's less value to acquiring it in advance.

The existence of shortages is a puzzle for pure classical economics; it requires that the price not increase to match quantity demanded with quantity supplied. Institutional factors of many kinds may prevent these adjustments — the employees of a price-offering retailer may not notice that the inventories are getting low until the shelves are empty, and those employees with the authority to raise prices may learn about it even later; the firm may have reputational reasons not to raise prices, if consumers are likely to punish the firm more for raising prices than for running out of stock; there may even be local regulations on prices, either hard caps or more ambiguous warnings about "price-gouging", that prevent the price from rising. Even if two consumer markets have the same instituional frictions on the supply-side, consumers may respond more forcefully to a prospective shortage in one market than another. Differences in both demand-side behavior and supply-side behavior further affect how agents form their expectations of the future, which then potentially feeds back into the behavior of both buyers and sellers. This paper seeks to address the question of what characteristics make a market more likely to get spooked and which characteristics make it more likely to recover. In some situations the answers to these questions may suggest policy responses, such as changing rationing rules imposed by policy-makers or adjusting formal price regulations; even where there are no policy responses, however, they are of interest to economic forecasters and anybody who wants to plan for the future.

Nichols and Zeckhauser (1977); Hendel, Dudine, and Lizzeri (2006); Bayer (2010); Mitraille and Thille (2014) study the stockpiling of storable goods in situations in which prices may differ from period to period. Their models have a small number of periods, strategic sellers, and markets that clear; this paper calls attention to an instability that exists in markets with inflexible prices and rationing, and which requires an infinite horizon.

Cavallo, Cavallo, and Rigobon (2014) study the retail market in Chile after its 2010 earthquake and Japan after its 2011 earthquake, where they note that prices changed very little even as shortages were severe and somewhat long-lasting — shortages in Chile were most severe a full two months after the initial disruption, and "a significant share of goods remained out of stock after six months." They especially note that "emergency" goods are more likely than others to stock-out after the earthquake, while "perishable" goods are less likely to do so. Roth (2008) identifies qualities that a market should have to be considered well-functioning; the one I am studying here he calls "thickness", but in financial literature especially it is "market liquidity". Several strains of theoretical literature are related to the emergence of market liquidity and its effect on the terms of trade of assets. Harrison and Kreps (1978) is seminal in a line of literature in which the ability of agents to trade an asset increases its value; in that model agents trade because they have different beliefs, and Morris (1996) characterizes the sets of beliefs they must hold to make the effect non-trivial. Other literatures depend on other asymmetries between agents to generate trade. Kiyotaki and Wright (1989, 1993) introduce a microfounded model of money by showing how the existence of bilateral markets with institutional barriers to multilateral trading can generate a liquidity premium for an asset. In those papers agents have different production capabilities and consumption preferences. Most of the financial economics papers on this topic use idiosyncratic "liquidity shocks" to motivate trading, as the agents suffering the shock seek to sell to the agents not suffering the shock. Rocheteau and Weill (2011) reviews a class of search models for assets in which the value of the asset is increased by the ease of selling it in the face of a liquidity shock. Duffie, Gârleanu, and Pedersen (2005, 2007) model a dealership-based market for financial markets and explore the attendant dynamics of price shocks. Vayanos and Wang (2011) provides a general three-period model (in the style of Diamond and Dybvig (1983)) to comprise several earlier similar models that explore the effects of different market imperfections on liquidity.

2 Model

My focus in this paper is the behavior of the demand side of the market, which will be assumed to have the convexity properties necessary for the separation theorems; accordingly, the demand side can be modeled separately from the supply side. The supply side, in fact, will not be modeled; the aims of the paper are intended not to be contingent on a particular model of the supply side of the market. Predictions for general market behavior will take supply behavior as input; policy recommendations will be made in terms of the supply behaviors that policymakers should attempt to cultivate.

2.1 Model

Each of a continuum of consumers receives an endowment of money m at each period t and uses it to buy b units of a storable consumption good; the consumer consumes c units of the good and $m_t - b_t p_t$ units of unmodeled goods, so that the utility in a given period is $u(c_t) - b_t p_t$. These other goods (and the quasilinear form of utility) keeps the marginal utility of money fixed and makes storing money undesirable.[3] While the consumption good can be stored, fraction λ of it depreciates in each period; $\lambda \to 1$ and $\lambda \to 0$ represent natural limits in which the good becomes nonstorable or nonperishable. Agents discount a period "flow" utility with a discount factor of δ. It is assumed that $\delta \in (0,1)$ and $\lambda \in [0,1]$. Letting C_t denote the amount of the good held at the beginning of the period, $b_t \geq 0$ the amount purchased in period t, and $c_t \in [0, C_t + b_t]$ the amount then consumed in period t, the agent's optimization problem is to maximize

$$\mathbb{E}_t \sum_{j=t}^{\infty} \delta^{j-t}(u(c_j) + m_j - b_j p_j)$$

subject to

$$\begin{aligned} C_{t+1} &= (1-\lambda)(C_t + b_t - c_t) \\ b_t &\geq 0 \\ c_t &\in [0, C_t + b_t] \end{aligned}$$

where each agent takes the price p_t as given. u is strictly increasing, strictly concave, and twice continuously differentiable.

As we move to general equilibrium, it will *not* be assumed that p_t is necessarily allowed to rise to clear the market. In those situations in which

[3]For example, it may be impossible to get a positive real return on stored money, whereupon the constant marginal utility of money and the discount factor make it suboptimal to store money.

p_t is held below the market-clearing price, there will be a shortage on the market; in the event of shortage, agents are assigned random priorities, based on which some agents get their chosen b_t, leaving the rest of the agents with $b_t = 0$.

The first order conditions for the optimizing consumer give

$$u'(c_t) \begin{cases} = p_t & b_t > 0 \\ \leq p_t & b_t = 0 \end{cases} \tag{1a}$$

$$u'(c_t) \begin{cases} \geq \delta(1-\lambda)\mathbb{E}_t\{u'(c_{t+1})\} & c_t = C_t + b_t \\ = \delta(1-\lambda)\mathbb{E}_t\{u'(c_{t+1})\} & c_t < C_t + b_t \end{cases} \tag{1b}$$

for agents choosing b_t and c_t to optimize expected discounted utility. (The b in (1a) is the agent's chosen b conditional on being able to purchase, while the b in (1b) is 0 if the agent was unable to buy.)

2.2 Equilibrium

We now consider the nature of equilibrium in this setting. We will consider "equilibrium" to be a sequence of prices, rationing levels, and quantities available for sale, starting at a particular time, along with a distribution of stockpiles with which agents enter the first period in which prices and rationing levels are given. Conditions for equilibrium will be that the quantities that optimizing agents buy in each period, with knowledge of current and future prices and rationing and current stockpiles, equal the quantity that is available for sale in that period.

A more natural definition of equilibrium might depend on specifying each agent's sequence of responses to that agent's history separately, and in using this simpler notion of "equilibrium" I am tacitly relying on a couple of properties of the model as I have stipulated it. One is that quantities and probabilities are specified exactly, such that the quantity demanded is equal to its expected value; one can think of a continuum of agents, and use the exact law of large numbers; alternatively, we could have a finite number of agents subject to the sequential service constraint of bank-run models, e.g. Wallace (1988), in which the macroeconomic quantity is exact and the probability results from a random ordering of which agents get to buy first. With this latter approach we still need to assume that the probability of being the last agent to buy and getting a "partially filled" order is 0, or at least that we are willing to neglect it. The other important tacit property is the serial independence of rationing; agents don't care much about their history beyond the level of stockpiles with which they enter a period; it doesn't affect future buying opportunities (conditional on the macroeconomic sequences).

If we are given an equilibrium starting at time t, then optimizing agents will buy and consume at time t, leading to a distribution of stockpiles entering time $t+1$; that distribution of stockpiles, along with the sequences of prices, rationing levels, and quantities available for sale starting at time $t+1$, will necessarily also be an equilibrium. We are not requiring the converse; an equilibrium beginning at time t need not be the natural consequence of any equilibrium beginning at time $t-1$. We are interested in evaluating when a (possibly measure-zero) "sunspot" event could cause agents to coordinate on a self-fulfilling macroeconomic course of action starting from time t, regardless of what expectations caused them to get to time t in a given state.

Formally, an equilibrium is

- A starting period $t \in \mathbb{Z}$
- A sequence of rationing levels $\{\pi_{t+k}\}_{k\geq 0}$ with $\pi_{t+k} \in [0,1]$

- A sequence of prices $\{p_{t+k}\}_{k\geq 0}$ with $p_{t+k} > 0$
- A sequence of quantities $\{q_{t+k}\}_{k\geq 0}$ with $q_{t+k} > 0$

such that

- Consumers choose functions for consumption $c : \mathbb{R} \to \mathbb{R}^+$ $(C_t \mapsto c_t)$ and purchases $b : \mathbb{R} \to \mathbb{R}^+$ $(C_t \mapsto b_t)$ to satisfy (1) for all times;
- $$\pi_t = \min\left\{1, \frac{q_t}{\langle b_t\rangle}\right\}$$
  for each t , where $\langle b_t\rangle$ represents an average over agents.

Note that with our strictly concave utility function, all consumers will have the same unique optimal strategy conditional on the macro variables; the equilibrium will be symmetric in this sense. Frequently we will specify

$$u = \frac{1}{1-\gamma}c^{1-\gamma} \qquad 0 < \gamma \neq 1$$

Note also that our attention is entirely on the demand side; there is no model of the supply side.

2.3 Optimizing Behavior

The approach to characterizing optimizing behavior by the consumer is as follows: in each period, the consumer will consume in such a quantity that the marginal utility of consumption is equal to the shadow value of the consumer's remaining stockpile. In a period in which the consumer is able to purchase, this shadow value is less than or equal to the price p_t at which that purchase would take place, with complementary slackness: the shadow value is equal to the price if the consumer purchases a positive quantity, and may only be less that the price if the consumer chooses not to purchase.

In a period in which the consumer is able to purchase the consumer will also purchase for *future* consumption if there is a possibility of rationing to prepare for. The evolution of shadow value μ when an optimizing agent is unable to buy is determined by a straightforward Euler equation, and the agent can therefore determine what the optimal conditional consumption in each future period would be as long as the agent is unable to purchase again between time t and that future period.[4] The agent will aim to end period t with enough of a stockpile to be able to consume that amount.

If each agent knows (and takes as given) the entire sequence of π_t and p_t, an optimizing agent's marginal utility of consumption will obey

$$\mu_t = \delta(1-\lambda)\left[\pi_{t+1}p_{t+1} + (1-\pi_{t+1})\mu_{t+1}\right]$$

where the μ_{t+1} on the right hand side is conditional on not being able to buy at time $t+1$; while the agent is unable to buy, μ will evolve according to

$$\mu_{t+1} = \frac{\delta^{-1}(1-\lambda)^{-1}}{(1-\pi_{t+1})}\mu_t - \frac{\pi_{t+1}}{(1-\pi_{t+1})}p_{t+1} \tag{2}$$

and μ_t will drop to p_t when the agent is able to buy.[5]

If we recursively define an effective discount factor

$$D_t = (1-\pi_t)\delta(1-\lambda)D_{t-1} \tag{3}$$

and if all $\mu_t \geq p_t$, then at time t an agent that last bought at time $t-j$ has

$$D_t\mu_{t,t-j} = D_{t-j}p_{t-j} - \sum_{i=0}^{j-1}\frac{\pi_{t-i}}{(1-\pi_{t-i})}D_{t-i}p_{t-i} \tag{4}$$

[4] There will be no equilibrium in which $\mu \leq 0$ at any time for any agent who was able to buy at some previous point; a positive mass of such agents would demand an unlimited quantity at that previous point, so quantity demanded would not be positive and finite.

[5] If μ can dip below p an agent would have to be allowed to sell some of the stockpile at the market price in order for these formulas to hold; note that positive quantities imply $\pi_{t+1} > 0$, and that future conditional values of μ are irrelevant if $\pi = 1$ between now and then, so we will frequently suppose that $\pi \in (0, 1)$.

and consumes $c_{t,t-j}$ such that $u'(c_{t,t-j}) = \mu_{t,t-j}$.

For u strictly concave, $\mu = u'(c)$ provides a one-to-one mapping between c and μ that holds at all times; this can be combined with the above relationships to determine what the agent's consumption pattern would look like if the agent never succeeds in buying again. If that is the case, then the agent will simply consume the amount in stock, adjusting for the deterioration of stocks over time:

$$(1-\lambda)^{-1}C_{t+1} = \sum_{j=1}^{\infty} \frac{c_{t+j,t}}{(1-\lambda)^j} \tag{5}$$

where $c_{t+j,t+1} = c_{t+j,t}$ if the agent doesn't buy at time $t+1$. If the agent does get to buy again, μ becomes p, and the sequence of c_t if the agent can't buy after *that* changes; an agent who does buy, having last bought at $t-k$, buys

$$\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(c^*_{t+j,t} - c^*_{t+j,t-k}\right)$$

(enough to match the new optimal consumption path). The fraction of agents in this group is

$$\pi_{t-k}(1-\pi_{t-k+1})\cdots(1-\pi_{t-1})\pi_t$$

so the total quantity purchased in a period is

$$\sum_{j=0}^{\infty}(1-\lambda)^{-j}c^*_{t+j,t}\pi_t - \sum_{k=1}^{\infty}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(c^*_{t+j,t-k}\right)\pi_{t-k}(1-\pi_{t-k+1})\cdots(1-\pi_{t-1})\pi_t \tag{6}$$

$$= \sum_{k=1}^{\infty}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(c^*_{t+j,t} - c^*_{t+j,t-k}\right)\pi_{t-k}(1-\pi_{t-k+1})\cdots(1-\pi_{t-1})\pi_t \tag{7}$$

When studying the recovery of a system from a shortage, it is useful to note that if $\pi_{t+j} = 1$ for some $j \geq 1$, then $c^*_{t+j,t-k} = 0$ for all $k \geq 0$, so the sum over j effectively gets cut off.

rationing:	0.5	0.875	0.9671	0.9737	0.9752	0.9763
prices:	1	1	1	1	1	1
quantities:	1	1	1	1	1	1
stockpile:	0.5	0.5543	0.5453	0.533	0.5207	0.5086
consumption:	0.5	0.9208	0.9814	0.985	0.9857	0.9862
stocktarget:	2	1.618	1.561	1.545	1.532	1.519
conditional	1	0.7324	0.173	0.02698	0.004039	0.00059
consumptions:		1	0.4829	0.08337	0.01251	0.001829
			1	0.4422	0.07272	0.01065
				1	0.4322	0.06911
					1	0.4239
						1

Figure 1: $\delta = 0.95$, $\lambda = 0.05$, $\gamma = 2$. The first 50% of the consumers that get to the store each buy twice as much as usual, so the store runs out; if they are unable to buy again in the first three periods, they consume 17.3% of the classical equilibrium amount in period 3.

For much of this paper, we will work with HARA utility, for which $\mu = c^{-\gamma}$ for some fixed γ, in which case

$$c_{t+j,t}^{-\gamma} = D_{t+j}^{-1} D_t \mu_t - D_{t+j}^{-1} \sum_{k=1}^{k=j} \frac{\pi_{t+k}}{(1 - \pi_{t+k})} D_{t+k} p_{t+k}$$

3 Markets with Multiple Equilibria

3.1 Illustrative Example

Consider the scenario outlined in figure 1, where agents know that prices will be constant, a such a level that if they could reliably buy at that price they would buy and consume (in each period) a quantity that is normalized

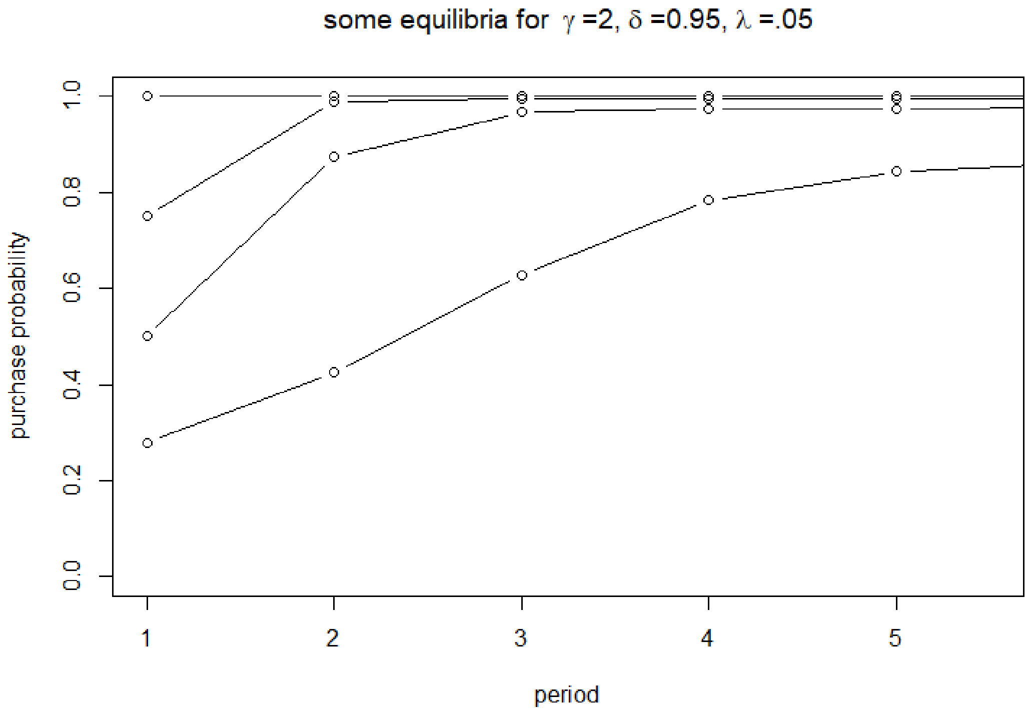

to 1. The agents also anticipate being unable to reliably buy, however; the probability of being able to buy in each of the first six periods is shown in the top row. The average agent enters period 2 with a quantity of .685 stored up from the previous period (some agents have 2.458 while others have 0); an agent able to buy in the second period, anticipating the rest of the sequence of probabilities, seeks to end the period with 2.037 in storage and to consume 1 in period 2, so the average agent wishes to buy 2.352; as only 42.5% of agents successfully buy, they buy a total quantity of $2.352 \times .4253 = 1$.

3.2 Main Result

Proposition 1. *Suppose*

- $u = (c^{(1-\gamma)})/(1-\gamma)$ *for some* $\gamma > 0, \gamma \neq 1$
- $\lambda > 0$

- *the quantity is positive, and is the same in every period*
- *the price is positive, and is the same in every period*
- *the price and quantity are such that* $\pi_{t+1} = \pi_{t+2} = \cdots = 1$ *completes an equilibrium*

then

1. *there are other equilibria with the same quantities and prices, but with* $\pi_{t+k} < 1 \forall k \geq 1$, *if* $\gamma > 1$
2. *there are no other equilibria with the same quantities and prices, but with* $\pi_{t+1} < 1$, *if* $\gamma < 1$.

We call the $\pi = 1$ equilibrium "the classical equilibrium".

If $\gamma < 1$, the classical equilibrium is "stable" in this sense; if $\gamma > 1$, however, the economy is at risk of a sunspot shortage — there will be patterns of behavior such that a shortage results entirely from self-fulfilling expectations.

3.2.1 Lambda

Note that with prices constant, an agent who is able to buy will consume such that the marginal cost of utility is equal to the price, while any agent who is unable to buy will have a strictly higher marginal cost of utility, which is to say a lower consumption. In any equilibrium as described, if there's rationing, then the agents are on average consuming less than they're buying.

One consequence of this is that rationing equilibria of the sort just described require $\lambda > 0$. If $\lambda = 0$, the aggregate stockpiles of the agents never decrease. With $\delta < 1$, in order for that to be consistent with optimizing behavior, the $\pi_t \leq \bar{\pi} < 1$ have to be bounded away from 1, which means that stokpiles are in fact growing without bound, which ultimately becomes inconsistent with its being optimal not to consume some of it.

3.3 Summary of Proof

Given $\gamma > 1$, there are values of $\pi_{t+1} < 1$ such that there is a sequence $\pi_{t+2}, \pi_{t+3}, \cdots$ such that total quantity demanded in each period is equal to the value it would obtain if all the π were 1.

The optimizing behavior as §2.3 defines a mapping from sequences of π to sequences of quantities demanded; each step $\pi \mapsto \mu \mapsto c^* \mapsto \tilde{q}$ is smooth. A compact (with the ℓ^p topology) subset K of the space of sequences $\pi_{t+2}, \pi_{t+3}, \cdots$ is constructed on which $\sum_{i=2}^{\infty}(\tilde{q}_i - 1)^2$ is defined and smooth, and attains its minimum on K only where $\tilde{q}_i = 1$ for all $i \geq 1$.

The bulk of the work is showing that $-\partial(\tilde{q}_1, \cdots, \tilde{q}_u)/\partial(\pi_{t+2}, \pi_{t+3}, \cdots, \pi_{t+u+1})$ is positive definite on the interior of K; some attention is also required to constructing K to ensure that the minimum of S on K is attained on the interior.

Details are relegated to the appendix (page 23).

3.4 Intuition of the Result

In this section we illustrate the dynamics of a market in a self-fulfilling shortage using approximations that hold for $\pi \approx 1$.

If $\epsilon = 1 - \pi_{t+1}$,

$$\mu_{t+1} = \frac{\delta^{-1}(1-\lambda)^{-1}\mu_t - p_{t+1}}{\epsilon} + p_{t+1}$$

If

$$\epsilon \ll \frac{\delta^{-1}\mu_t}{p_{t+1}} - 1,$$

then
$$\mu_{t+1} \approx \frac{\delta^{-1}(1-\lambda)^{-1}\mu_t}{\epsilon}$$
If prices are constant and the probability of rationing is much less than $\delta^{-1}-1$, then the top-up level is

$$\approx p^{-1/\gamma} + \delta^{1/\gamma}(1-\lambda)^{1/\gamma-1}\epsilon_{t+1}^{1/\gamma}p^{-1/\gamma}$$

(but somewhat higher for positive ϵ) with purchases

$$\approx p^{-1/\gamma}(1-\epsilon_t)\left(1+\delta^{1/\gamma}(1-\lambda)^{1/\gamma-1}\left(\epsilon_{t+1}^{1/\gamma}-(1-\lambda)\epsilon_t^{1/\gamma}\right)\right)$$

$\approx p^{-1/\gamma}+$

$$p^{-1/\gamma}\left(\delta^{1/\gamma}(1-\lambda)^{1/\gamma-1}\left(\epsilon_{t+1}^{1/\gamma}-(1-\lambda)\epsilon_t^{1/\gamma}\right)-\epsilon_t\right)$$

The quantities purchased are more or less constant and equal to $p^{-1/\gamma}$ if

$$(1-\lambda)^{-1}\epsilon_{t+1}^{1/\gamma} = \delta^{-1/\gamma}(1-\lambda)^{-1/\gamma}\epsilon_t + \epsilon_t^{1/\gamma} \tag{8}$$

which gives
$$\frac{\epsilon_{t+1}}{\epsilon_t} = \left(\delta^{-1/\gamma}(1-\lambda)^{1-1/\gamma}\epsilon_t^{1-1/\gamma}+(1-\lambda)\right)^{\gamma}$$
For $\gamma > 1$ and $\lambda > 0$,[6] as $\epsilon_t \to 0$,

$$\frac{\epsilon_{t+1}}{\epsilon_t} = (1-\lambda)^{\gamma}$$

while for $\gamma < 1$, if ϵ_t is small, $\epsilon_{t+1} > \epsilon_t$; in that case there is no asymptotically non-rationing equilibrium with positive rationing. (If ϵ_t is too small, the right hand side exceeds 1.)

In words, an anticipated shortage at time $t+1$ results in higher quantity demanded at time t, creating a shortage at time t; if $\gamma > 1$ the shortage created at time t is larger than the shortage anticipated at the next period. If

[6]For $\lambda = 0$, the expression stays bigger than 1 for small positive ϵ.

$\gamma < 1$, then a shortage at time t will only occur in response to the expectation of ever more severe shortages. Within a finite number of periods, there is no $\pi_{t+k+1} \in [0, 1]$ that would cause each agent to demand a high enough quantity to support the preceding sequence of πs.

3.5 Welfare analysis

Consider agents with ex ante total utility of

$$\mathbb{E}_t \sum_{j=t}^{\infty} \delta^{j-t}(u(c_j) - b_j p_j)$$

with u strictly concave, increasing, and differentiable in c. u is a strictly convex function of μ (e.g. from (Rockafellar, 1970, theorem 23.5)); for example, in the CRRA framework in which we're mostly working, current consumption utility is

$$u = \frac{1}{1-\gamma} c^{1-\gamma} = \frac{1}{1-\gamma} \mu^{1-1/\gamma}$$

Proposition 2. *Consider two equilibria with*

- *the same quantities*
- *the same prices*
- $\pi = 1$ *at all times in one of the equilibria.*

The the equilibrium with $\pi = 1$ *Pareto dominates that with* $\pi < 1$.

The sums

$$\mathbb{E}_t \sum_{j=t}^{\infty} \delta^{j-t} u_j \text{ and } \mathbb{E}_t \sum_{j=t}^{\infty} \delta^{j-t} b_j p_j$$

converge separately. As long as $\mu_t \geq p_t$, $\mathbb{E}_t \sum_{j=t}^{\infty} \delta^{j-t} u_j$ is lower than it would be in the classical equilibrium; on the other hand, as long as agents start out symmetrically, $\mathbb{E}_t \sum_{j=t}^{\infty} \delta^{j-t} b_j p_j$ is the same for any two equilibria in which agents buy (in aggregate) the same quantity (matching period-by-period).

For fixed sequences of p and π,

$$T_t = \sum_{j=0}^{\infty} \frac{c_{t+j,t}}{(1-\lambda)^j} (5)$$

is strictly convex and strictly decreasing in μ_t — each term in the sum is — thus μ_t is stictly convex and strictly decreasing as a function of T_t

As long as $\delta^{-1}(1-\lambda)^{-1}\mu_t > p_{t+1}$ and $\delta^{-1}(1-\lambda)^{-1}p_{t+k} > p_{t+k+1}$ for all k, T_t is monotonically decreasing[7] in future π, for any fixed μ_t; the entire curve shifts down if any π increases toward 1. Thus μ_t as a function of T_t also shifts down (because μ is strictly decreasing in T); at least for small enough stockpiles that $\mu_t > \delta(1-\lambda)p_{t+1}$, agents coming in with stockpiles will also have a reduction in welfare from shortages.

4 Supply response: eliminating rationing

The result above requires that prices be unresponsive to shortages. If prices rise, then quantity demanded goes down; if buyers anticipate that prices *would* rise sufficiently if there *were* a shortage, then they also anticipate that the shortage will cease if it begins, and the bad equilibria are eliminated. To be clear, this is ultimately off-path; prices don't actually change, but it is commonly believed that they would if there were rationing.[8]

[7]Strictly decreasing in π_{t+l} if $\pi_{t+k} < 1 \forall k : 1 \leq k < l$; otherwise it is unchanged.

[8]This simplifies the welfare analysis by obviating the concern that higher prices hurt consumers, and need to be balanced against reliability. This hypothetical welfare concern does, however, introduce a question of policy credibility: if there is a policy that allows prices to rise a lot, it may be harder to establish the common belief that it would be followed. It is worth noting, therefore, that the shortage in figure 1 can be avoided by threatening to raise the price by 6%.

Our primary interest is in $q = 1$ at all times, so “price flexibility” will not be p_{t+k} as a function of q_{t+k} as it typically is. Instead, we allow p_{t+k} to be a function of the sequence of π. Given a sequence of such functions (one function for each k), a sequence of π implies a sequence of p. Our equilibrium concept thus changes somewhat for this section: instead of defining the equilibrium as sequences of π, q, and real numbers p that satisfy demand-side constraints, an equilibrium is a sequence of π, q, and functions p that, evaluated on the sequence π, give a sequence of real numbers p that satisfy the requirements of equilibrium from §2.2. In exploring the structure of equilibria, we still think of q and p as given and seek sequences of π that satisfy the criteria jointly with them.

There are a lot of sequences of functions p that we could explore, even restricting ourselves to those that are continuously differentiable and for which there is an equilibrium with $\pi_{t+k} = 1$ and $q_{t+k} = 1$ for all k. (We will restrict ourselves to that class of functions.) A formally simple subclass of such functions is those that depend only on rationing in the previous period, $p_{t+k}(\pi_{t+k-1})$, with $p(1) = 1$ and p decreasing in π_{t+k-1} (i.e. increasing as π decreases). An economically more intuitive simple subclass is that for which each p_{t+k} is a function of C_{t+k}, with $p(0) = 1$ and p increasing in C_{t+k}.

As before, the existence of equilibria in which $\pi_{t+k} < 1$ but $\pi_{t+k} \to 1$ as $k \to \infty$ will hinge on the behavior of

$$\frac{\partial q_{t+k+j}}{\partial \pi_{t+k}}$$

for different values of j when the various π are near 1. Because the total quantity demanded $\tilde{q}$ is smooth in p and π and p is continuously differentiable in π,

$$\frac{\partial \tilde{q}_t}{\partial \pi_{t+k}} = \left.\frac{\partial \tilde{q}_t}{\partial \pi_{t+k}}\right|_p + \sum_m \left.\frac{\partial \tilde{q}_t}{\partial p_{t+m}}\right|_\pi \frac{\partial p_{t+m}}{\partial \pi_{t+k}}$$

as long as the sum is well-defined (its terms absolute summable). If the sum is much smaller (in absolute size) than $\partial\tilde{q}_t/\partial\pi_{t+1}$, then it will not affect whether there is such an equilibrium; price-sensitivity that is big enough to matter has to diverge as $\pi \to 1$ at least as quickly as $\partial\tilde{q}_t/\partial\pi_{t+1}$ and $\partial C_{t+1}/\partial\pi_{t+1}$ do. We therefore concentrate on p as a function of C.

4.1 Price Response to Stockpiles

If $\epsilon = 1 - \pi_{t+1}$,

$$\mu_{t+1} = \frac{\delta^{-1}(1-\lambda)^{-1}\mu_t - p_{t+1}}{\epsilon} + p_{t+1}$$

If

$$\epsilon \ll \frac{\delta^{-1}\mu_t}{p_{t+1}} - 1,$$

then

$$\mu_{t+1} \approx \frac{\delta^{-1}(1-\lambda)^{-1}\mu_t}{\epsilon}$$

If prices are constant and the probability of rationing is much less than $\delta^{-1} - 1$, then the top-up level is

$$\approx p^{-1/\gamma} + \delta^{1/\gamma}(1-\lambda)^{1/\gamma-1}\epsilon_{t+1}^{1/\gamma}p^{-1/\gamma}$$

with purchases

$$\approx (1-\epsilon_t)\left(p^{-1/\gamma}\left(1 + \delta^{1/\gamma}(1-\lambda)^{1/\gamma-1}\epsilon_{t+1}^{1/\gamma}\right) - C_t\right)$$

Suppose p is a function of C; for present purposes, $p = 1 + \alpha C$. Now the quantity purchased q_t obeys

$$q_t + \epsilon_t q_t \approx 1 - \frac{\alpha}{\gamma}C_t + \delta^{1/\gamma}(1-\lambda)^{1/\gamma-1}\epsilon_{t+1}^{1/\gamma} - C_t$$

with $q_t = 1$ if

$$\epsilon_t \approx \delta^{1/\gamma}(1-\lambda)^{1/\gamma-1}\epsilon_{t+1}^{1/\gamma} - \frac{\alpha+\gamma}{\gamma}C_t$$

Now, suppose again that this is all anticipated by optimizing agents, and that $C_t \approx \delta^{1/\gamma}(1-\lambda)^{1/\gamma}\epsilon_t^{1/\gamma}$, with $\gamma > 1$ such that (with $\epsilon \approx 0$) $\epsilon \ll \epsilon^{1/\gamma}$, then

$$\epsilon_{t+1}^{1/\gamma} \approx \frac{\alpha+\gamma}{\gamma}(1-\lambda)\epsilon_t^{1/\gamma},$$

such that there is, as for $\gamma < 1$, no asymptotically vanishing rationing equilibrium if

$$\frac{\alpha+\gamma}{\gamma}(1-\lambda) > 1,$$

which is to say if

$$\alpha > \left((1-\lambda)^{-1} - 1\right)\gamma \tag{9}$$

If agents anticipate that (for whatever reason) prices would rise sufficiently quickly if stockpiles rise, then these equilibria go away.

5 Modeling questions

5.1 Interpretting γ

We are working with a fairly specific functional form for utility: additively separable between money and the consumption good, and linear in money, such that the marginal utility of wealth is constant. The results generalize somewhat, and in intuitive ways.

First, note that if the marginal utility of wealth is η, $u'(c) \leq \eta p$ with equality for $c > 0$; we've fixed $\eta = 1$ on the grounds that it is unlikely to change very much in a shortage, but the γ parameter may be easier to interpret in terms of how consumption changes with income. If η is allowed to vary then γ is proportional to the elasticity of consumption with respect to the marginal utility of wealth, to wit

$$\begin{aligned} c^* &= \arg\max_c m + u(c) - \eta(m + c/p + I) \Rightarrow \\ u'(c^*) &= \eta/p \\ \frac{\partial u'(c^*)}{\partial c^*} &= \frac{1}{p}\frac{\partial \eta}{\partial c^*} \end{aligned}$$

if u is twice differentiable and increasing at c^*;

$$\begin{aligned} \frac{c^*}{u'(c^*)}\frac{\partial u'(c^*)}{\partial c^*} &= \frac{c^*}{u'(c^*)}\frac{1}{p}\frac{\partial \eta}{\partial c^*} \\ &= \frac{c^*}{\eta}\frac{\partial \eta}{\partial c^*}. \end{aligned}$$

For a given income elasticity of utility, then, γ is proportional to the income elasticity of consumption for an optimizing agent. Accordingly we will identify low γ with "luxury" goods and high γ with "necessities".

A second point to note is that the result is driven by the relationship between c and u at low values of c; if the utility of consumption is such that

$$\frac{cu''}{u'}$$

isn't constant but is bounded above or below 1 as $c \to 0$ then our results should generalize to that utility function.

6 Conclusion

I have noted that, in markets with fixed prices and, where necessary, rationing, there can be shortages even when the quantity supplied would be enough to satisfy contemporaneous demand. I have demonstrated that this holds for storable goods if they are particularly "inelastic" in that consumers' willingness to pay rises very steeply as their consumption goes down, and that some price flexibility can prevent this.

This is useful for private agents wishing to predict which markets are susceptible to such "runs"; it also suggests some policy responses. Most clearly, this suggests an additional cost to any policies that are likely to reduce price flexibility, and an additional benefit to policies that reduce stickiness, especially in the face of shortages. Quantity flexibility could replace price flexibility; this result argues on the side of retaining a national strategic reserve of petroleum and petroleum products and signalling a willingness to release them in response to temporary shortages.

I have noted previously that Cavallo, Cavallo, and Rigobon (2014) find empirical results in the face of actual shortages that are similar to the results of this model. They studied supply disruptions, however; there is no "classical equilibrium" with non-shortage prices in that case. After many disasters, however, the retail-level shortage outlasts the restoration of pre-disaster supplies; in those situations, even where policy-makers might wish to mitigate supply disruptions in other ways, any steps that encourage prices to adjust in response to on-going retail disruptions once the supply disruptions are over would help stabilize the market more quickly.

7 Appendix

7.1 c and its Derivatives

The quantity purchased at time t has been written as a function of π and c^*; the c^* have been written as functions of μ. Here we place some bounds on the derivative of the q with respect to the π, holding the p constant, by use of the chain rule; first we take the derivative of μ, then of c^*, and then q.

From (4),

$$\mu_{t,t-j} = D_t^{-1} D_{t-j} p_{t-j} - \sum_{i=0}^{j-1} \frac{\pi_{t-i}}{(1-\pi_{t-i})} D_t^{-1} D_{t-i} p_{t-i};$$

$$\frac{\partial (D_t^{-1} D_{t-j})}{\partial \pi_{t-k}} = \frac{D_t^{-1} D_{t-j}}{1-\pi_{t-k}} \text{ if } 0 \le k < j$$

$$\frac{\partial (D_t^{-1} D_{t-i} \pi_{t-i}/(1-\pi_{t-i}))}{\partial \pi_{t-i}} = \frac{D_t^{-1} D_{t-i}}{(1-\pi_{t-i})^2}$$

If $0 \le k < j$

$$\begin{aligned}\frac{\partial \mu_{t,t-j}}{\partial \pi_{t-k}} =& \frac{D_t^{-1} D_{t-j}}{1-\pi_{t-k}} p_{t-j} - \frac{D_t^{-1} D_{t-k}}{(1-\pi_{t-k})^2} p_{t-k} \\ & - \sum_{i=k+1}^{j-1} \frac{D_t^{-1} D_{t-i}}{1-\pi_{t-k}} \frac{\pi_{t-i}}{(1-\pi_{t-i})} p_{t-i}\end{aligned}$$

while otherwise it's 0.

Note that if, for some fixed $\bar{\epsilon} > 0$, $1 - \pi_{t+k} \le \bar{\epsilon}$ for $i \le k \le j$, then

$$\frac{\partial \mu_{t,t-j}}{\partial \pi_{t-k}} \le \bar{\epsilon}^{j-1} \delta^j (1-\lambda)^j p_{t-j}.$$

Of particular interest, given μ_{t-1}, for any t, we have

$$\frac{\partial \mu_t}{\partial \pi_t} = \frac{\delta^{-1}(1-\lambda)^{-1}\mu_{t-1} - p_t}{(1-\pi_t)^2}$$

Now we can find the derivative of c^* using the chain rule;

$$\frac{\partial c^*_{t+k,t}}{\partial x} = \frac{\partial \mu_{t+k,t}^{-1/\gamma}}{\partial x} = -\frac{1}{\gamma} \mu_{t+k,t}^{-1/\gamma - 1} \frac{\partial \mu_{t+k,t}}{\partial x}$$

$$= -\frac{1}{\gamma} \frac{c^*_{t+k,t}}{\mu_{t+k,t}} \frac{\partial \mu_{t+k,t}}{\partial x}$$

If, for some fixed $\bar{\epsilon} > 0$, $1 - \pi_{t+k} \le \bar{\epsilon}$ for $i \le k \le j$, then

$$\left| \frac{\partial c^*_{t,t-j}}{\partial \pi_{t-k}} \right| = \left| -\frac{1}{\gamma} \mu_{t,t-j}^{-1/\gamma-1} \frac{\partial \mu_{t,t-j}}{\partial \pi_{t-k}} \right| \le \frac{1}{\gamma} \mu_{t,t-j}^{-1/\gamma-1} \bar{\epsilon}^{j-1} \delta^j (1-\lambda)^j p_{t-j}. \tag{10}$$

Using (4) again,

$$\mu_{t,t-j} = D_t^{-1} D_{t-j} p_{t-j} - \frac{\pi_{t-(j-1)}}{(1-\pi_{t-(j-1)})} D_t^{-1} D_{t-(j-1)} p_{t-(j-1)} - \sum_{i=0}^{j-2} \frac{\pi_{t-i}}{(1-\pi_{t-i})} D_t^{-1} D_{t-i} p_{t-i}$$

For $\bar{\epsilon} < 1/2$ and $j \geq 2$, we can write this as

$$D_t^{-1} D_{t-j} \left(p_{t-j} - \delta(1-\lambda)\pi_{t-(j-1)} p_{t-(j-1)} - \sum_{i=0}^{j-2} \frac{\pi_{t-i}}{(1-\pi_{t-i})} D_{t-j}^{-1} D_{t-i} p_{t-i} \right)$$
$$\geq D_t^{-1} D_{t-j} p \left(1 - \delta(1-\lambda)(1-\bar{\epsilon}) - \frac{\bar{\epsilon}(1-\bar{\epsilon})}{1-\delta(1-\lambda)\bar{\epsilon}} \right)$$

when p is constant. Thus for $j \geq 2$,

$$\left| \frac{\partial c^*_{t,t-j}}{\partial \pi_{t-k}} \right| \leq \frac{1}{\gamma} \bar{\epsilon}^{j-1} \delta^j (1-\lambda)^j p^{-1/\gamma} \delta^{j/\gamma+j} (1-\lambda)^{j/\gamma+j} \bar{\epsilon}^{j/\gamma+j} \left(1 - \delta(1-\lambda)(1-\bar{\epsilon}) - \frac{\bar{\epsilon}(1-\bar{\epsilon})}{1-\delta(1-\lambda)\bar{\epsilon}} \right)^{-1/\gamma-1}$$

$$\left| \frac{\partial c^*_{t,t-j}}{\partial \pi_{t-k}} \right| \leq \bar{\epsilon}^{j/\gamma+2j-1} \frac{p^{-1/\gamma}}{\gamma} \left(\delta(1-\lambda) \right)^{j/\gamma+2j} \left(1 - \delta(1-\lambda) + \delta(1-\lambda)\bar{\epsilon} - \frac{\bar{\epsilon}(1-\bar{\epsilon})}{1-\delta(1-\lambda)\bar{\epsilon}} \right)^{-1/\gamma-1} \tag{11}$$

Given μ_{t-1},

$$\frac{\partial c_{t,t-k}}{\partial \pi_t} = -\frac{1}{\gamma} \mu_{t,t-k}^{-1/\gamma-1} \frac{\partial \mu_t}{\partial \pi_t} = -\frac{1}{\gamma} \mu_{t,t-k}^{-1/\gamma-1} \frac{\delta^{-1}(1-\lambda)^{-1}\mu_{t-1} - p_t}{(1-\pi_t)^2}$$
$$= -\frac{1}{\gamma} \mu_{t,t-k}^{-1/\gamma-1} \frac{\mu_t - p_t}{1-\pi_t}$$

7.2 Constructing a Domain

We have (2)

$$\mu_{t+1} = \frac{\delta^{-1}(1-\lambda)^{-1}}{(1-\pi_{t+1})} \mu_t - \frac{\pi_{t+1}}{(1-\pi_{t+1})} p_{t+1};$$

in particular,

$$\mu_{t+1,t} = \frac{\delta^{-1}(1-\lambda)^{-1}}{(1-\pi_{t+1})} p_t - \frac{\pi_{t+1}}{(1-\pi_{t+1})} p_{t+1}$$
$$\mu_{t,t-1} = \frac{\delta^{-1}(1-\lambda)^{-1}}{(1-\pi_t)} p_{t-1} - \frac{\pi_t}{(1-\pi_t)} p_t$$
$$\mu_{t+1,t-1} = \frac{\delta^{-1}(1-\lambda)^{-1}}{(1-\pi_{t+1})} \left(\frac{\delta^{-1}(1-\lambda)^{-1}}{(1-\pi_t)} p_{t-1} - \frac{\pi_t}{(1-\pi_t)} p_t \right) - \frac{\pi_{t+1}}{(1-\pi_{t+1})} p_{t+1}$$

For $p_{t+1} = p_t = p_{t-1} = 1$, $\gamma > 1$, and (given) $\pi_t < 1$ but sufficiently close to 1,

$$\mu_{t,t-1}^{-1/\gamma} + (1-\lambda)^{-1}\mu_{t+1,t-1}^{-1/\gamma} - (1-\lambda)^{-1}\mu_{t+1,t}^{-1/\gamma} = 1/\pi_t - 1$$

has a solution for $\pi_{t+1} \in (1-(1-\lambda)(1-\pi_t), 1)$. Let $\bar{\epsilon}_t = 1-\pi_t$, then define $\bar{\epsilon}_{t+1}$ as $1-$ that π_{t+1} that solves this equation; we similarly define $\bar{\epsilon}_{t+k+1}$ in terms of $\bar{\epsilon}_{t+k}$, and thus define a sequence of $\bar{\epsilon}$. We will define

$$K = \prod_k [0, \bar{\epsilon}_{t+k}] \subset \ell^2 \tag{12}$$

Now, consider

$$1 - c^*_{t,t-1} + (1-\lambda)^{-1}\left(c^*_{t+1,t} - c^*_{t+1,t-1}\right)$$

which $= 1/(1-\bar{\epsilon}_t) - 1$ if $\pi_t = 1-\bar{\epsilon}_t$ and $\pi_{t+1} = 1-\bar{\epsilon}_{t+1}$. It is decreasing in π_{t+1} and is increasing in π_t; thus, for any $\pi_t \in [1-\bar{\epsilon}_t, 1)$, if $\pi_{t+1} = 1-\bar{\epsilon}_{t+1}$,

$$\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(c^*_{t+j,t} - c^*_{t+j,t-1}\right) \geq 1 - c^*_{t,t-1} + (1-\lambda)^{-1}\left(c^*_{t+1,t} - c^*_{t+1,t-1}\right) \geq 1/(1-\bar{\epsilon}_t) - 1$$

(6) gives $q_t =$

$$\sum_{j=0}^{\infty}(1-\lambda)^{-j}c^*_{t+j,t}\pi_t + \sum_{k=1}^{\infty}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(-c^*_{t+j,t-k}\right)\pi_{t-k}(1-\pi_{t-k+1})\cdots(1-\pi_{t-1})\pi_t$$

and since

$$\sum_{k=1}^{\infty}\left(c^*_{t+j,t-k}\right)\pi_{t-k}(1-\pi_{t-k+1})\cdots(1-\pi_{t-1}) \leq c^*_{t+j,t-1}$$

we therefore have the result that, if $\pi_t \in [1-\bar{\epsilon}_t, 1]$ and $\pi_{t+1} = 1-\bar{\epsilon}_{t+1}$, demanded $q_t \geq 1$. On the other hand, if $\pi_{t+1} = 1$, then $c_{t+j,t} = 0$ for $j \geq 1$, and $q_t =$

$$\pi_t - \sum_{k=1}^{\infty} c^*_{t,t-k}\pi_{t-k}(1-\pi_{t-k+1})\cdots(1-\pi_{t-1})\pi_t \leq 1$$

with equality only if $\pi_t = 1$ as well.

Note that K defined in (12) is compact in ℓ^2; in particular, it is totally bounded, as, given any $\delta > 0$, $\exists N : \sum_{k=N}^{\infty}\bar{\epsilon}^2_{t+k} < \delta^2/4$, and $\prod_{k<N}[0, \bar{\epsilon}_{t+k}]$ can be covered with a finite number of balls of radius $\delta\sqrt{3}/2$.

7.3 derivatives of q

(6) gives $q_t =$

$$\sum_{j=0}^{\infty}(1-\lambda)^{-j}c^*_{t+j,t}\pi_t+\sum_{k=1}^{\infty}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(-c^*_{t+j,t-k}\right)\pi_{t-k}(1-\pi_{t-k+1})\cdots(1-\pi_{t-1})\pi_t$$

$\partial q_t/\partial\pi_{t+k} =$

$$1_{t+k=t}\sum_{j=0}^{\infty}(1-\lambda)^{-j}c^*_{t+j,t}+\pi_t\sum_{j=0}^{\infty}(1-\lambda)^{-j}\frac{\partial c^*_{t+j,t}}{\partial\pi_{t+k}}$$
$$-\sum_{l=1}^{\infty}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\frac{\partial c^*_{t+j,t-l}}{\partial\pi_{t+k}}\pi_{t-l}(1-\pi_{t-l+1})\cdots(1-\pi_{t-1})\pi_t$$
$$+\sum_{l=1}^{\infty}\left(\frac{1_{t-l=t+k}}{\pi_{t-l}}+\frac{1_{t=t+k}}{\pi_t}\right)\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(-c^*_{t+j,t-l}\right)\pi_{t-l}(1-\pi_{t-l+1})\cdots(1-\pi_{t-1})\pi_t$$
$$+\sum_{l=1}^{\infty}1_{t-l+1\le t+k\le t-1}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(c^*_{t+j,t-l}\right)\pi_{t-l}(1-\pi_{t-l+1})\cdots(1-\pi_{t-1})\pi_t/(1-\pi_{t+k})$$

$\partial q_t/\partial\pi_{t+k} =$

$$1_{t+k=t}\sum_{j=0}^{\infty}(1-\lambda)^{-j}c^*_{t+j,t}+\pi_t\sum_{j=0}^{\infty}(1-\lambda)^{-j}\frac{\partial c^*_{t+j,t}}{\partial\pi_{t+k}}$$
$$+1_{t=t+k}\sum_{l=1}^{\infty}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(-c^*_{t+j,t-l}\right)\pi_{t-l}(1-\pi_{t-l+1})\cdots(1-\pi_{t-1})$$
$$-\sum_{l=\max\{1,1-k\}}^{\infty}\sum_{j=\max\{0,k\}}^{\infty}(1-\lambda)^{-j}\frac{\partial c^*_{t+j,t-l}}{\partial\pi_{t+k}}\pi_{t-l}(1-\pi_{t-l+1})\cdots(1-\pi_{t-1})\pi_t$$
$$+1_{t+k\le t-1}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(-c^*_{t+j,t+k}\right)(1-\pi_{t+k+1})\cdots(1-\pi_{t-1})\pi_t$$
$$+1_{t+k\le t-1}\sum_{l=1-k}^{\infty}\sum_{j=0}^{\infty}(1-\lambda)^{-j}\left(c^*_{t+j,t-l}\right)\pi_{t-l}(1-\pi_{t-l+1})\cdots(1-\pi_{t-1})\pi_t/(1-\pi_{t+k})$$

For small $\bar{\epsilon}$, if $t = t+k$ this is dominated by the $c^*_{t,t}$ term; the difference between $\partial q_t/\partial\pi_t$ and $c^*_{t,t}$ can be bounded by positive powers of $1-\pi_t$. Similarly, if $t+k > t$, the dominating term is $\pi_t(1-\lambda)^{-k}\frac{\partial c^*_{t+k,t}}{\partial\pi_{t+k}}$, and for $t+k < t$, we have primarily $-c^*_{t,t+k}(1-\pi_{t+k+1})\cdots(1-\pi_{t-1})\pi_t$. The diagonal elements of

$$\frac{\partial(q_1,\cdots,q_u)}{\partial(\pi_2,\cdots,\pi_{u+1})}$$

are $\pi_t(1-\lambda)^{-1}\frac{\partial c^*_{t+1,t}}{\partial\pi_{t+1}}$, which is bounded above and below by positive multiples of $-(1-\lambda)^{-1}\bar\epsilon_{t+1}^{1/\gamma-1}$ (as shown in the previous section); the terms above the diagonal are bounded (in absolute value) by $(1-\lambda)^{-1}\bar\epsilon_{t+1}^{1/\gamma}\bar\epsilon_{t+2}^{1/\gamma}\cdots\bar\epsilon_{k}^{1/\gamma-1}$. The terms immediately below the diagonal correspond to $\partial q_t/\partial\pi_t$, and are bounded above and below by positive constants; the terms below that are bounded in absolute value by a positive multiple of $\bar\epsilon_{t-k+1}^{1+1/\gamma}\cdots\bar\epsilon_{t-1}^{1+1/\gamma}\bar\epsilon_{t}^{1/\gamma}$. A crude sketch is provided in figure 2.

For $\gamma > 1$, the terms on the diagonal get arbitrarily far from zero for small $\bar\epsilon$, (to $-\infty$), while all other terms stay bounded. Consider matrix M in which each row of

$$\frac{\partial(q_1,\cdots,q_u)}{\partial(\pi_2,\cdots,\pi_{u+1})}$$

has been divided by the diagonal term; matrix M has (by construction) 1's along the diagonal and, for sufficiently small starting $\bar\epsilon$,[9] the sum of the squares of the off-diagonal terms add up to less than 1. M is therefore positive definite;[10] therefore

$$-\frac{\partial(q_1,\cdots,q_u)}{\partial(\pi_2,\cdots,\pi_{u+1})} \text{ is positive definite for } \gamma > 1$$

7.4 Synthesis

Let $S=\sum_i(q_i-1)^2$, which we regard as a function of π; note, from arguments in §7.2, that the terms in the sum are bounded by multiples of positive powers

[9]Because the $\bar\epsilon_t$ converge exponentially to 0, a fixed $\bar\epsilon$ can be chosen for all u, viz. any size of matrix we consider.

[10]See §7.5.

$$\frac{\partial(q_1,\cdots,q_u)}{\partial(\pi_2,\cdots,\pi_{u+1})} = \begin{pmatrix} -\bar\epsilon^{1/\gamma-1} & -\bar\epsilon^{2/\gamma-1} & -\bar\epsilon^{3/\gamma-1} & \cdots \\ \bar\epsilon^{0} & -\bar\epsilon^{1/\gamma-1} & -\bar\epsilon^{2/\gamma-1} & \cdots \\ -\bar\epsilon^{1/\gamma} & \bar\epsilon^{0} & -\bar\epsilon^{1/\gamma-1} & \cdots \\ -\bar\epsilon^{1+2/\gamma} & -\bar\epsilon^{1/\gamma} & \bar\epsilon^{0} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

Figure 2: A heuristic look at the Jacobian. This is a charicature, but provides a good starting point for intuition. The diagonal term is bounded away from zero; $\bar\epsilon_t$ decreases with t, so the power counting should be regarded as a rough guide, even as a bound.

of $\bar\epsilon_i$ and thus S converges on the domain K. K is compact, and S is smooth, and S thus attains it minimum on K.

The gradient is $\nabla S = 2\sum_i (q_i-1)\nabla(q_i)$; because the (negative) Jacobian $\partial q/\partial\pi$ is positive-definite, $\sum_j (q_j-1)\nabla_j S < 0$ unless $(q_i-1)=0$ for all i. Similarly, anywhere on the boundary of K, if $\hat n$ is any vector normal to a binding constraint and pointed into K, $\hat n\cdot\nabla S<0$; thus the minimum of S must occur in the interior of K, where $q_t=1$ for all t.

7.5 Positive Definiteness of Almost Diagonal Matrices

Suppose $v\in\mathbb{R}^n$ and i an integer between 1 and n inclusive, and that $\exists j\neq i$: $v_j\neq 0$. Then

$$\min_{m\in\mathbb{R}^n} m^Tm \text{ s.t. } m_i=1, m^Tv=0$$

is given by a unique m, for which $m_j=-v_iv_j/C$ for $j\neq i$, where $C=\sum_{j\neq i} v_j^2$; the value of the minimum is $m^Tm=$

$$1+\frac{v_i^2}{v^Tv-v_i^2}$$

Now consider $v \in \mathbb{R}^n$, $v \neq 0$, and we wish to solve

$$\min_{M_{n\times n}} \sum_{i=1}^{n}\sum_{j=1}^{n} M_{ij}^2 \text{ s.t. } Mv = 0_{n\times 1}, M_{ii} = 1 \forall i$$

If only one element of v is non-zero, the constraints cannot be satisfied; otherwise, the problem separates along rows of M, and we can apply the previous result to each row of M, giving

$$\sum_i \left(1 + \frac{v_i^2}{v^T v - v_i^2}\right) > n + \sum_i \frac{v_i^2}{v^T v} = n + 1.$$

The result that interests us is the contrapositive: given $M_{n\times n}$ such that $M_{ii} = 1\forall i$ and $\sum_{ij:i\neq j} M_{ij}^2 \leq 1$, $\nexists v \in \mathbb{R}^n, v \neq 0 : Mv = 0_{n\times 1}$. The immediate corollary is

Proposition 3. *Given $M_{n\times n}$ such that $M_{ii} = 1\forall i$ and $\sum_{ij:i\neq j} M_{ij}^2 \leq 1$, M is non-singular.*

Note that if a matrix satisfies the conditions of the proposition, then so does its symmetric part; furthermore, the set $\mathcal{S}$ of symmetric matrices satisfying those conditions is connected in the usual Euclidean topology in $\mathbb{R}^{n^2}$. Finally, $1_n \in \mathcal{S}$, thus the eigenvalues of any element of $\mathcal{S}$ are positive, and thus

Proposition 4. *Given $M_{n\times n}$ such that $M_{ii} = 1\forall i$ and $\sum_{ij:i\neq j} M_{ij}^2 \leq 1$, M is positive definite.*

8 Bibliography

Bayer, R.-C. (2010): "Intertemporal price discrimination and competition," *Journal of Economic Behavior & Organization*, 73(2), 273–293.

Cavallo, A., E. Cavallo, and R. Rigobon (2014): "Prices and Supply Disruptions during Natural Disasters," *Review of Income and Wealth*, 60(S2), S449–S471.

Diamond, D. W., and P. H. Dybvig (1983): "Bank runs, deposit insurance, and liquidity," *Journal of Political Economy*, 91(3), 401–419.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005): "Over-the-counter Markets," *Econometrica*, 73(6), 1815–1847.

——— (2007): "Valuation in Over-the-Counter Markets," *Review of Financial Studies*, 20(6), 1865–1900.

Harrison, J. M., and D. M. Kreps (1978): "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *The Quarterly Journal of Economics*, 92(2), 323–36.

Hendel, I., P. Dudine, and A. Lizzeri (2006): "Storable Good Monopoly: The Role of Commitment," *American Economic Review*, 96(5), 1706–1719.

Kiyotaki, N., and R. Wright (1989): "On Money as a Medium of Exchange," *Journal of Political Economy*, 97(4), 927–54.

——— (1993): "A Search-Theoretic Approach to Monetary Economics," *American Economic Review*, 83(1), 63–77.

Mitraille, S., and H. Thille (2014): "Speculative storage in imperfectly competitive markets," *International Journal of Industrial Organization*, 35(C), 44–59.

Morris, S. (1996): "Speculative Investor Behavior and Learning," *The Quarterly Journal of Economics*, 111(4), 1111–1133.

Nichols, A. L., and R. J. Zeckhauser (1977): "Stockpiling Strategies and Cartel Prices," *Bell Journal of Economics*, 8(1), 66–96.

Rocheteau, G., and P.-O. Weill (2011): "Liquidity in Frictional Asset Markets," *Journal of Money, Credit and Banking*, 43, 261–282.

Rockafellar, R. T. (1970): *Convex Analysis*. Princeton University Press.

Roth, A. E. (2008): "What Have We Learned from Market Design?," *Economic Journal*, 118(527), 285–310.

Vayanos, D., and J. Wang (2011): "Theories of Liquidity," *Foundations and Trends in Finance*, 6(4), 221–317.

Wallace, N. (1988): "Another attempt to explain an illiquid banking system: the Diamond and Dybvig model with sequential service taken seriously," *Quarterly Review*, (Fall), 3–16.