

Online appendix for “Price Pressures”

Terrence Hendershott and Albert J. Menkveld

This document has the following supplemental material:

1. Section 1 presents the infinite horizon version of the Ho and Stoll (1981) model and solves it numerically. Fig. 7 graphs its solution.
2. Table 5 extends Table 3 by providing all parameters instead of just a subset.

1. A recursive version of the Ho and Stoll (1981) model and its numerical solution

1.1. Recursive dynamic inventory model setup

Ho and Stoll (1981) are the first to set up the dynamic program for the risk-averse intermediary and cast it in a finite time horizon where utility is generated only by final wealth. To better fit our analysis of liquidity supplier inventory and prices on the NYSE, we take the model, cast it in recursive form with an infinite horizon, and solve for the intermediary's optimal policy and the resulting stationary distribution.¹

1.1.1. Liquidity demand process

Ho and Stoll (1981) use a reduced form model for the transaction rates (where transaction size is fixed at one unit) of public buyers and sellers that was first proposed by Garman (1976). These transaction rates are linear in the ask and bid price, with the public buy rate decreasing in the ask price and public sell rate increasing in the bid price.

The social cost analysis requires a bit more structure on the liquidity demand process. Instead of separating the two sides of the market, we generalize the investor arrival process by considering agents who arrive at the market and depending on the prices quoted decided whether to buy, sell, or not trade. The investors incur idiosyncratic endowment shocks that can be hedged by an offsetting position in a (single) risky security that trades in the market. This private value v is stochastic and zero in expectation as it is a private value relative to a security's common value that we fix at zero throughout for ease of notation.² The agent

¹ Others also study dynamic inventory control. Madhavan and Smidt (1993) further enriches the Ho and Stoll (1981) setting by adding an informed liquidity demander. Amihud and Mendelson (1980) analyze a dynamic inventory model in recursive form where a risk-neutral intermediary maximizes average profit and stays away from an exogenously given maximum inventory constraint as arriving at that state limits her profit opportunity (she can essentially exit on one side). We work from the Ho and Stoll (1981) model as we focus on the price risk associated with an suboptimal inventory position.

² The common value here is the equilibrium price in the absence of the liquidity friction that is modeled. It is the price at which all idiosyncratic shocks wash in the crosssection of investors. We refer to the investors' idiosyncratic desire to trade as their hedging values, although their private motivations for trade could be more general.

arrives at the market and trades only if v is above the ask (agent buys a unit) or below the bid (agent sells a unit).³ We parameterize this process as follows. The idiosyncratic shocks happen, and agents arrive at the market according to a Poisson process with rate $2\lambda\theta$. Conditional on arrival, the private value v is uniformly distributed with support $[-\frac{1}{\theta}, \frac{1}{\theta}]$. This specification implies Fig. 6 for the public buy (sell) rate at the ask (bid) price, where the arrival rate per private value is derived as the arrival rate ($2\lambda\theta$) times the private value density ($\frac{1}{2}\theta$) and the public buy and sell rates follow by integrating over private value rates from the ask to $\frac{1}{\theta}$ and from $\frac{1}{\theta}$ to the bid, respectively. The total private value rate (or the size of the market) is obtained by setting the bid and ask price equal to zero and is thus represented by the size of the triangle defined by $-\frac{1}{\theta}$ and $\frac{1}{\theta}$ on the price axis and $\lambda\theta$ on the transaction rate axis. The size of this area is λ , which represents the size of maximum consumer surplus. The realized private value rate depends on the bid and ask price set by the intermediary who produces liquidity. If she is risk-neutral and, therefore, operates on a zero cost rate, she charges the monopoly price $p = \frac{1}{2\theta}$ and, for public buys, captures a producer surplus rate equal to the square defined by A in the graph. The consumer surplus rate on this side of the market is the triangle B, and the deadweight loss due to monopoly power is C on this side of the market.

[Insert Figure 6 near here]

The second parameter θ captures the dispersion of private values and, therefore, governs the price sensitivity of the net transaction rate (defined as the public buy rate minus the public sell rate) standardized by the investor arrival rate. Lowering of both the bid quote and the ask quote by s increases the rate of buys minus sells by $2s\lambda\theta^2$, which, standardized by the arrival rate of $2\lambda\theta$, equals θs . Comparative statics show that this is an important parameter in the intermediary's optimization behavior. θ parallels the price elasticity of demand that a monopolist faces in a standard product market.

1.1.2. *The intermediary's dynamic program*

We set up the intermediary's optimization in discrete time to solve for her optimal policy and obtain a stationary distribution of pricing errors and inventory. This distribution characterizes the series' time series properties, i.e., contemporaneous as well as lagged moments. The Poisson private demand arrivals are modeled as discrete probabilities that are linear in prices.⁴ The intermediary is risk-averse and maximizes

³ If agents have the opportunity to contact the market more than once then v can be interpreted as the hedging benefit from the current trading opportunity until the investor's next contact with the market [see Weill (2007) and Lagos et al. (2011) for more on investors' intertemporal optimization].

⁴ This is a reasonable approximation as the Poisson process itself is the limiting case of event probabilities that are linear in small discrete time increments.

expected utility where her utility is the standard time-separable CRRA utility. She is a monopolist liquidity supplier who produces liquidity supply by issuing firm price quotes that are take-it-or-leave-it prices that liquidity demanders see when arriving at the market. She can therefore not price discriminate and capture the full consumer surplus.

Let $V(\cdot)$ be the maximum utility that the intermediary can achieve when starting off with an inventory of i_0 shares and a wealth of w_0 :

$$V(i_0, w_0) = \max_{\{c_t, a_t \leq \frac{1}{\theta}, b_t \geq -\frac{1}{\theta}\}_{t=0}^{\infty}} E[\sum_{t=0}^{\infty} \beta^t u(c_t) | i_0, w_0] \quad (38)$$

subject to the budget constraint

$$w_{t+1} = R(w_t - c_t) + i_t \Delta m_{t+1} + a_t \max(-\Delta i_t, 0) - b_t \max(\Delta i_t, 0) \quad (39)$$

and subject to the following transition probabilities across inventory states:

$$P[i_{t+1} = i_t - 1 | i_t] = \int_{a_t}^{\frac{1}{\theta}} (2\lambda\theta) \left(\frac{1}{2}\theta\right) dv = \lambda\theta(1 - \theta a_t) \quad (40)$$

and

$$P[i_{t+1} = i_t + 1 | i_t] = \lambda\theta(1 + \theta b_t),$$

where i_t is the intermediary's inventory position at time t , w_t is her wealth, c_t is her consumption at time t , β is her discount factor, R is the gross risk-free return, Δm_t is her stochastic dividend in period $t+1$ (which runs from time t to time $t+1$), which has zero expectation and a variance equal to σ^2 , and b_t (a_t) is the bid (ask) price she sets at time t . The stochastic dividend stream is a modeling device to minimize the accounting in the model. We emphasize that no sources of uncertainty are ignored (see also Ho and Stoll, 1981, p. 52).

We exploit the recursive nature of the problem that leads to the following Bellman equation for the

inventory state i (using the law of iterated expectations):

$$\begin{aligned}
V(i, w) = & \max_{c_{iw}, a_{iw}, b_{iw}} u(c_{iw}) + \beta E_{\Delta m} \left[V(i+1, R(w - c_{iw}) + i\Delta m - b_{iw}) \right] \lambda \theta (1 + \theta b_{iw}) + \\
& \beta E_{\Delta m} \left[V(i-1, R(w - c_{iw}) + i\Delta m + a_{iw}) \right] \lambda \theta (1 - \theta a_{iw}) + \\
& \beta E_{\Delta m} \left[V(i, R(w - c_{iw}) + i\Delta m) \right] (1 - \lambda \theta (1 + \theta b_{iw}) - \lambda \theta (1 - \theta a_{iw})),
\end{aligned} \tag{41}$$

where Δm is the stochastic dividend in the oncoming period.

With upper bounds on the absolute inventory position and wealth it is straightforward to show that the functional equations defined by Eq. (41) define a contraction in the (i, w) space by verifying Blackwell's pair of sufficient conditions (see Ljungqvist and Sargent, 2004, p. 1012). This implies the existence of a unique fixed point in the space of bounded continuous functions and therefore guarantees the existence of a unique equilibrium.

The numerical solution illustrates that the intermediary seems to endogenously stay away from large inventory and wealth upper bounds for our calibration in which relative risk aversion is larger than one and the reciprocal of the discount rate (β^{-1}) is strictly larger than the risk-free rate (R). The endogenous upper bound on wealth is best understood based on Huggett (1993), which shows that for exogenous nondiversifiable endowment risk the two parameter restrictions guarantee that the agent does not let wealth grow infinitely large. The intuition is that with CRRA utility the agent's willingness to bear absolute risk grows with wealth. The reciprocal of the discount rate being strictly larger than the risk-free rate makes her prefer consuming today over tomorrow. It is, therefore, perfectly intuitive that at some level of wealth she prefers consuming out of wealth over accumulating more given that risk is less of a consideration on high wealth levels.

The model we propose extends this setting as the exposure to idiosyncratic risk (inventory) is decided upon endogenously. The same intuition applies for why she does not let wealth grow to infinity. In addition, the intermediary does not visit extreme inventory states as her earning potential does not change with inventory, whereas her cost does given the stochastic dividend exposure ($i\Delta m$). In the numerical solution we set the bounds wide enough so that the intermediary decides never to visit the upper bound states.

1.2. *The numerical solution: the optimal policy and implied stationary distribution*

We numerically solve the intermediary’s dynamic program, which is summarized by Eq. (38). The calibration follows, to the extent possible, the base case parametrization proposed by (Ho and Stoll, 1981, p. 67). We present the calibration details in Section 2.

Fig. 7 presents the numerical solution to the intermediary’s dynamic program. Panel A plots the value function which is defined as the maximum discounted utility the intermediary obtains conditional on starting off in a particular inventory-wealth state. Its concavity is generated by the time discounting in the dynamic program (see Ljungqvist and Sargent, 2004, Section A.2.). Given the value function and the Bellman equations, it is a straightforward one-period optimization to establish the intermediary’s optimal control policy (i.e., bid price, ask price, and consumption) in each inventory-wealth state. This control policy, in turn, determines the system transition laws. If she is started off with sufficient wealth and zero inventory, she eventually ends up in the stationary distribution over the inventory-wealth states, which is also depicted in Panel A. The stationary distribution illustrates that she endogenously chooses to stay far from the inventory and wealth upper and lower bounds.

[Insert Figure 7 near here]

Panel B of Fig. 7 illustrates the optimal control policy, which consists of the bid price, the ask price, and the consumption in each of the inventory-wealth states. The graphs provide the midquote and the bid-ask spread instead of the bid and ask price separately to better illustrate price pressures. The graph confirms the intuition that the intermediary skews the midquote to generate an transactions that mean-reverts her inventory. She lowers the midquote on long positions and raises it on short positions.

Two observations are particularly useful to motivate the econometric model we propose for the empirical analysis. First, the intermediary’s policy seems to support the first-order approximation we use in the econometric model, which is that price pressure is linear in inventory and does not depend on wealth. The wealth result is particularly useful given that we do not have access to data on the intermediary’s wealth level. Second, relative to midquote changes the bid-ask spread appears constant across states.⁵ It is only slightly higher than the spread that a risk-neutral intermediary would charge: $2(2\theta)^{-1} = 0.2$.⁶ This is a

⁵ Zabel (1981) and Mildenstein and Schleef (1983) find that the spread is independent of inventory. Ho and Stoll (1981) find inventory has a very small effect on the spread.

⁶ It is not surprising that the price of immediacy, i.e., the spread, is slightly higher because the marginal cost of production is no longer zero but becomes positive when the intermediary is risk-averse.

reassuring result as it confirms the Grossman and Miller (1988) argument that the spread captures the cost of processing orders (or monopoly rents in our case), whereas midquote changes capture the dynamic trade-off the intermediary does between staying one more period on a costly suboptimal position or discounting prices to mean-revert that position. This apparent orthogality of monopoly rents that are reflected in the spread and limited risk-bearing capacity that lead to price pressures suggests that competition in the intermediation sector on unchanged overall risk-bearing capacity drives down the bid-ask spread to marginal cost but leaves price pressures unchanged. That is, the fundamental trade-off of discounting price versus bearing idiosyncratic risk for an additional period remains. This is our intuition and to formally derive a general result is beyond the scope of the current paper.

The consumption graph illustrates that the intermediary insures against adverse dividend shocks by maintaining a wealth buffer. She consumes less than expected earnings on low wealth levels and more than expected earnings on high wealth levels. This is best illustrated by the zero inventory policies as her expected earnings of a potential current period transaction equal a constant $2 \times (1 - 0.5 \times 0.23) \times 0 = 0.98$ as prices are unchanged across the wealth dimension. She consumes less than these expected earnings on low wealth levels and thus insures against adverse future dividend shocks through savings. She consumes more than expected earnings and thus eats out of her large wealth buffer on high wealth levels. This consumption pattern makes wealth mean-reverting. Furthermore, she reduces her consumption when on a larger nonzero inventory, which reflects the higher need for a wealth buffer given that she is in a more risky state.

2. Details on the calibration and the solution method

2.1. Base case model calibration

We numerically solve the intermediary's dynamic program, which is summarized by Eq. (38). The calibration follows, to the extent possible, the base case parametrization proposed by (Ho and Stoll, 1981, p. 67).

1. The intermediary's coefficient of relative risk aversion (ρ) is two.
2. The time interval length is one day.
3. The arrival rate function is $g(x) = \lambda_0(1 - \theta x)$, where $\lambda_0 = 1$ and $\theta = 5$ so that each day the representative investor arrives and considers a trade. The transaction size is ten units.⁷

⁷ Ho and Stoll (1981) use the same θ but set $\lambda_0 = 2,000$ as the annual transaction rate. To match their arrival rate we set

4. The stochastic dividend risk on one inventory unit equals $\sigma = \sqrt{0.50/200} = 0.05$, where one year contains two hundred trading days.
5. The intermediary's discount rate β is 0.9 to capture the two week horizon Stoll proposes for the intermediary.⁸
6. The intermediary earns a gross daily risk-free rate of $R = 1 + 0.10/200$ on her savings for the next period.

We then discretize the set of inventory-wealth states (i, w) by choosing a grid with carefully selected bounds for both the inventory and the wealth dimension of the state. We choose upper bounds for inventory and wealth so that the state-space becomes a finite set. We set inventory bounds equal to -50 and 50 and verify that ex post these inventory states appear to be nonbinding constraints.⁹ The step size for the inventory dimension is governed by the transaction size ten. We choose a wealth upper bound of ten where we again verify that the limit does not bind endogenously. We set the wealth lower bound equal to zero, which is an absorbing bankruptcy state from which point on consumption is arbitrarily small.¹⁰ The wealth step size is constrained by the stochastic dividend risk, which is implemented as a discrete random variable where the outcome is plus or minus σ with equal probability. As the inventory dimension contains only multiples of ten, the stochastic dividend commands a (maximum) wealth step size of $10 \times 0.05 = 0.5$.¹¹

We then set up the discretized version of the Bellman equation, which is a trivial extension of the system defined by Eq. (41). The step size on the wealth dimension necessarily discretizes the bid and ask prices that the intermediary can choose, but we can get as close as one might desire to a continuum of prices by shrinking the wealth step size at the cost of computational speed. After discretizing, we solve the dynamic program numerically by iterating on the Bellman equation, relying on the contraction property of the recursive mapping to achieve convergence.

the transaction size equal to ten with the interpretation that one successful daily transaction means $2,000/200 = 10$ intraday transactions of one inventory unit.

⁸ A 0.9 discount rate implies that roughly two-thirds of the total discounted value of a dividend stream of one falls into the first two weeks.

⁹ One has to be careful here, as in addition to it being a very risky state, the intermediary internalizes that she can exit this state only on one side (i.e., less inventory), which limits her earning potential. Ideally, we would like to let the inventory upper bound I tend to infinity to rule out this nonrisk explanation. We simulate such limit numerically and verify that the solution does not change by taking larger upper bounds.

¹⁰ CRRA utility satisfies the Inada conditions so that intermediaries steer clear of this state.

¹¹ Technically, the intermediary could still end up strictly outside of wealth interval if she gets a large enough draw on the stochastic dividend. In the implementation, we steer her to the lower or upper bound of wealth in the cases that this happens. This does not affect our solution as the intermediary endogenously chooses to stay away from those states.

2.2. Bellman equations for the interior of the inventory-wealth state domain

We arrive at the following discrete state dynamic program:

$$\begin{aligned}
V(i, w) &= \max_{\tilde{w}_{iw}, a_{iw}, b_{iw}} u(w - R^{-1}\tilde{w}_{iw}) & (42) \\
&+ \beta E_{\Delta m} [V(i, \tilde{w} + i\Delta m)](1 - g(a_{iw}) - g(b_{iw})) \\
&+ \beta E_{\Delta m} [V(i + 1, \tilde{w} + i\Delta m + b_{iw}\Delta i) + \Delta m]g(b_{iw}) \\
&+ \beta E_{\Delta m} [V(i - 1, \tilde{w} + i\Delta m - a_{iw}\Delta i) + \Delta m]g(a_{iw}) \\
g(x) &= 1 - \theta x \quad (\text{probability of a transaction}) \\
i &\in \{-50, -40, \dots, 50\} \quad w \in \{0, 0.5, \dots, 10\},
\end{aligned}$$

where the three expected value terms correspond to no-arrival, seller-arrival, and buyer-arrival, respectively. Eq. (42) defines the Bellman equations for the nonboundary inventory-wealth states. But, before turning to the boundary Bellman equations, we need to make sure that we stay on the inventory-wealth grid in the iterations. This motivates a binomial distribution for stochastic dividend

$$P[\Delta m = -\sigma] = P[\Delta m = +\sigma] = 0.5 \quad (43)$$

and restricts the set of admissible controls to

$$a_{iw}, b_{iw} \in \{0, 0.05, 0.10, \dots\} \text{ s.t. } g(a_{iw}), g(b_{iw}), 1 - g(a_{iw}) - g(b_{iw}) \geq 0 \quad (44)$$

$$\tilde{w}_{iw} \in \{0, 0.5, \dots, \min(10, 0.5[2Rw])\} \quad (45)$$

The upper bound for end-of-period wealth \tilde{w}_{iw} prevents negative current period consumption.

2.3. Bellman equations for the boundary of the inventory-wealth state domain

As for the boundary inventory states, we adjust the Bellman equation defined in Eq. (42) by reducing the set of admissible controls. For the maximum inventory state, for example, we restrict the bid price to $b_{Iw} = \theta^{-1}$ (which effectively sets the seller arrival rate to zero). As for the boundary wealth states, in the zero wealth state we keep the intermediary on a low enough consumption level relative to her earning power in the nonzero wealth state to make her endogenously choose to stay away from the bankruptcy state. In states

close to the maximum wealth state, we redirect her to the maximum wealth state if she enjoys a stochastic dividend that would make her transit to a larger wealth than the maximum wealth. One interpretation is that the government taxes excessive wealth away in these high wealth states. We reiterate that these assumptions on the maximum inventory or maximum wealth Bellman equations are inconsequential as we choose the maximum inventory and maximum wealth boundary large enough so that the intermediary endogenously chooses never to get near to these states (see discussion in footnote 11).

From this point on, we follow (Ljungqvist and Sargent, 2004, Subsection 4.2, p. 95), i.e., we iterate on the Bellman equation and the contraction property of the recursive mapping guarantees convergence to the unique solution.

Table 5

State space model estimates.

This table reports the parameter panels that were omitted from Table 3.

Quintile	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	All	<i>t</i> -statistic
<i>Panel I: $\sigma(u)_t$ error term permanent volatility (in basis points)</i>														
Q1	116	117	121	126	164	187	241	183	177	120	101	98	146	
Q2	133	127	130	134	171	201	251	200	196	141	118	116	160	
Q3	140	133	136	144	185	200	241	196	193	141	128	127	164	
Q4	151	144	144	145	182	199	224	197	185	148	133	140	166	
Q5	166	162	168	167	211	222	246	225	225	182	161	174	192	
All	141	136	140	143	182	202	240	200	195	146	128	131	166	
<i>Panel J: $\sigma(f)_t$ common factor volatility (in basis points)</i>														
All	57	46	67	98	122	106	152	130	154	98	70	63	97	
<i>Panel K: β_i</i>														
Q1	0.99	0.98	1.04	0.99	0.88	0.73	0.32	0.79	0.99	0.97	0.90	0.94	0.88	14.46
Q2	1.09	0.87	0.95	0.85	0.82	0.55	0.35	0.85	1.00	1.08	1.04	1.11	0.88	12.51
Q3	1.15	1.02	0.93	0.82	1.02	0.56	0.48	0.84	1.06	1.11	1.21	1.14	0.94	13.21
Q4	1.12	0.72	0.67	0.73	0.90	0.48	0.35	0.77	0.90	1.01	1.25	1.24	0.84	10.41
Q5	1.02	0.69	0.58	0.63	0.87	0.36	0.43	0.79	0.77	0.93	1.28	1.20	0.79	9.90
All	1.07	0.86	0.83	0.80	0.90	0.54	0.39	0.81	0.94	1.02	1.13	1.13	0.87	13.66
<i>Panel L: β_i^0</i>														
Q1	0.07	0.01	0.03	0.05	0.06	0.09	0.18	-0.04	-0.02	0.01	0.01	0.02	0.04	2.27
Q2	-0.11	0.01	-0.09	-0.03	0.01	0.02	0.13	-0.12	-0.08	-0.09	-0.04	-0.02	-0.03	-1.62
Q3	-0.22	-0.15	-0.12	-0.12	-0.16	-0.06	0.08	-0.08	-0.13	-0.15	-0.11	0.02	-0.10	-4.18
Q4	-0.33	-0.12	-0.10	-0.16	-0.20	-0.08	0.14	-0.07	-0.03	-0.08	0.01	0.10	-0.08	-2.19
Q5	-0.35	-0.19	-0.08	-0.18	-0.32	-0.09	-0.12	-0.27	-0.12	-0.20	-0.15	0.05	-0.17	-5.31
All	-0.19	-0.09	-0.07	-0.09	-0.12	-0.02	0.08	-0.12	-0.08	-0.10	-0.06	0.03	-0.07	-3.38
<i>Panel M: β_i^1</i>														
Q1	0.04	0.01	0.02	0.03	0.06	0.03	0.09	-0.04	0.01	-0.02	0.02	-0.02	0.02	1.81
Q2	-0.02	0.05	-0.03	-0.00	0.04	-0.03	0.08	-0.09	0.00	-0.04	-0.00	-0.00	-0.00	-0.29
Q3	-0.09	-0.03	-0.04	-0.02	-0.04	-0.07	0.06	-0.05	-0.07	-0.06	-0.07	0.08	-0.03	-2.27
Q4	-0.13	-0.01	-0.01	-0.04	-0.13	-0.04	0.08	-0.09	-0.08	-0.05	-0.07	0.13	-0.04	-1.64
Q5	-0.20	-0.07	0.02	-0.06	-0.19	-0.04	-0.03	-0.18	-0.15	-0.13	-0.14	0.12	-0.09	-3.29
All	-0.08	-0.01	-0.01	-0.02	-0.05	-0.03	0.06	-0.09	-0.06	-0.06	-0.05	0.06	-0.03	-2.11
<i>Panel N: β_i^2</i>														
Q1	0.04	-0.01	0.02	0.03	0.04	0.04	0.04	-0.01	-0.01	-0.01	-0.00	-0.04	0.01	1.70
Q2	-0.03	0.03	-0.03	-0.01	0.04	-0.03	0.05	-0.08	0.01	-0.03	-0.01	0.01	-0.01	-0.74
Q3	-0.04	-0.02	-0.07	-0.02	-0.03	-0.06	-0.00	-0.07	-0.02	-0.04	-0.02	0.07	-0.03	-2.70
Q4	-0.06	0.02	-0.06	-0.07	-0.09	-0.07	0.03	-0.11	0.01	0.00	-0.04	0.11	-0.03	-1.66
Q5	-0.14	-0.01	-0.03	-0.06	-0.09	-0.07	-0.03	-0.15	-0.04	-0.03	-0.10	0.12	-0.05	-2.72
All	-0.05	0.00	-0.03	-0.03	-0.03	-0.04	0.02	-0.09	-0.01	-0.02	-0.03	0.05	-0.02	-2.21
<i>Panel O: β_i^3</i>														
Q1	-0.00	0.01	0.00	0.00	0.01	0.01	0.04	-0.00	0.01	-0.00	-0.00	-0.02	0.00	1.02
Q2	-0.03	0.03	0.00	0.00	0.03	-0.01	0.04	-0.03	0.01	-0.03	-0.01	-0.02	-0.00	-0.27
Q3	0.00	0.02	-0.01	0.00	-0.01	-0.02	0.00	-0.02	-0.00	-0.02	-0.04	0.03	-0.01	-1.12
Q4	-0.03	0.05	-0.03	-0.02	-0.02	-0.02	0.04	-0.03	0.03	0.00	-0.04	0.08	-0.00	-0.03
Q5	-0.05	0.02	0.00	-0.02	-0.03	-0.02	0.00	-0.06	-0.00	-0.01	-0.03	0.08	-0.01	-1.00
All	-0.02	0.03	-0.01	-0.01	-0.00	-0.01	0.02	-0.03	0.01	-0.01	-0.02	0.03	-0.00	-0.50





Fig. 6. Consumer and producer surplus.

Panel A: Value function and stationary distribution



Panel B: Optimal controls: midquote, bid-ask spread, and consumption

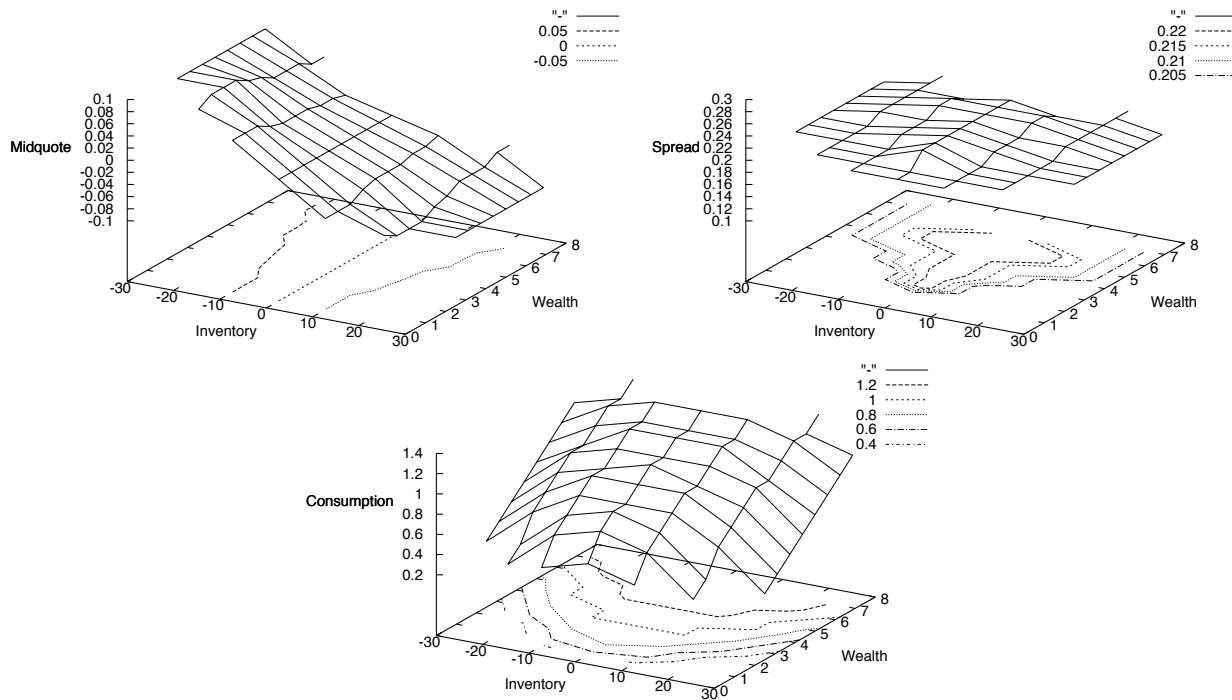


Fig. 7. Value function, stationary distribution, and optimal control. This figure illustrates the numerical solution to a dynamic programming problem of an intermediary who maximizes expected utility while quoting bid and ask prices to liquidity demanders whose arrival rates are less than perfectly price elastic. Panel A plots the value function (left) and the stationary distribution (right) over all the inventory-wealth states that the intermediary could find herself in. Panel B illustrates her optimal control by plotting the midquote (left), the bid-ask spread (right), and her consumption (bottom) conditional on her inventory-wealth state. The parametrization of the problem is standard and follows Ho and Stoll (1981). The intermediary has constant relative risk aversion utility with a coefficient of relative risk aversion equal to two. She operates at a daily frequency and has a two-week horizon that we capture by a (daily) discount factor of 0.9. Without loss of generality, we fix the fundamental price at zero and capture inventory price risk through a stochastic dividend with mean zero and daily standard deviation of \$0.05. The probability of a public buyer (seller) arrival depends linearly on the ask (bid) price, where the slope θ captures the elasticity of arrival rates to prices (the benchmark is no friction, i.e., $\theta = \infty$).