Online Appendix for

Identification and Estimation of Forward-looking Behavior: The Case of Consumer Stockpiling

Andrew Ching and Matthew Osborne

10 Appendix: Figures

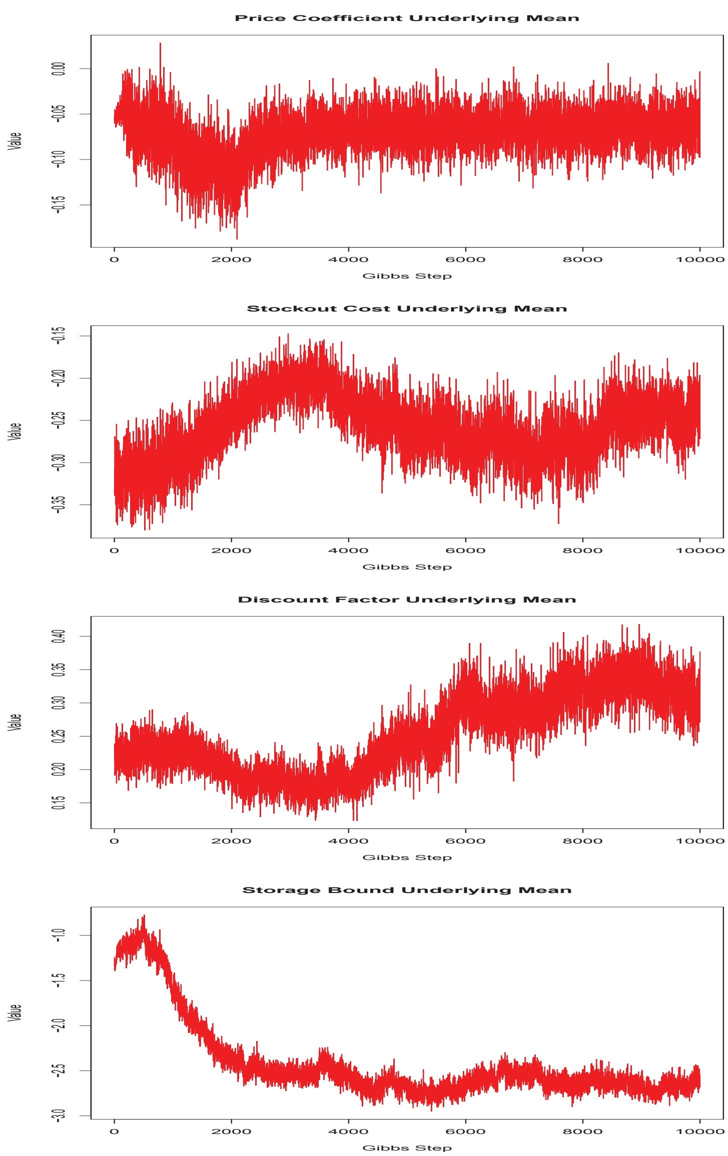

Figure 9: Appendix Figure: Plots of Gibbs Draws for Selected Dynamic Parameters

11 Appendix: Extension of Theoretical Model to Stochastic Consumption Rate

This section shows how the proofs in the paper of Propositions 1 and 2 can be extended to a model with stochastic consumption shocks. In this extension, we assume that the consumption shock takes on two values, 1 and 2, and that the probability a low value shock is π_c. We first prove the following Lemma which corresponds to Lemma 1 from the paper:

Lemma 4 *The value function $V(I)$ is increasing in inventory for $I < b+3$ and sufficiently small ω_1.*

Proof.

To begin, we propose an optimal policy which is that consumers only purchase when they run out. Under this optimal policy the value functions will be:

$$V(0) = -p + \pi_c \beta V(b-1) + (1-\pi_c)\beta V(b-2) - \omega_1,$$

$$V(1) = \pi_c \beta V(0) + (1-\pi_c)(-p + \beta V(b-1) - \omega_1),$$

$$V(2) = \pi_c(\beta V(1) - \omega_1) + (1-\pi_c)\beta V(0), ...$$

$$V(I) = \pi_c \beta V(I-1) + (1-\pi_c)\beta V(I-2) - \omega_B, ...$$

where B is the number of bottles at the end of the period.

We've defined the value functions as a system of equations we can solve. In particular for any inventory level I we can solve backwards until we've expressed $V(I)$ in terms of $V(1)$ and $V(0)$. Let's write

$$V(I) = A_I V(1) + B_I V(0) - C_I \omega_1.$$

For $I \leq b+1$. For $I > b+2$ there will be terms involving ω_2, etc, but we won't worry about these for now. The terms A_I, B_I and C_I are positive and are functions of β and π_c. We'll focus more on these terms later but for now one useful thing to note is that

$$A_I = \pi_c \beta A_{I-1} + (1-\pi_c)\beta A_{I-2},$$

and the same formula holds for B_I. For $C(I)$ the relevant equation is $C_I = \pi_c \beta C_{I-1} + (1 - \pi_c)\beta C_{I-2} + 1$.

Notice that if we do this we can solve for $V(1)$ and $V(0)$ as follows:

$$V(0) = \frac{((1-\pi_c)(\beta A_{N-1} - A_N) - 1)p + (C_N[(1-\pi_c)\beta A_{N-1} - 1] - A_N(1-\pi_c)[\beta C_{N-1} + 1])\omega_1}{(B_N - 1)((1-\pi_c)\beta A_{N-1} - 1) - A_N(\pi_c\beta + (1-\pi_c)B_{N-1})}$$

Now for $V(1)$ we get

$$V(1) = \frac{[(1-\pi_c)(B_n - \beta B_{N-1}) + \pi_c(1-\beta) - 1]p + [(B_N - 1)(1-\pi_c)(\beta C_{N-1} + 1) - C_N(\pi_c\beta + (1-\pi_c)\beta B_{N-1})]\omega}{(B_N - 1)((1-\pi_c)\beta A_{N-1} - 1) - A_N(\pi_c\beta + (1-\pi_c)B_{N-1})}$$

In the steps below, we will show the value function is increasing when $\omega_1 = 0$. Then, we will use continuity of the value function in ω_1 to argue it will still be increasing for small positive ω_1 values.

Next we show that $V(0) < 0$, $V(1) > V(0)$, and $V(2) > V(1)$. Let's start on $V(0)$. First we will show that the numerator $(1-\pi_c)(\beta A_{N-1} - A_N) - 1 < 0$. Note that we can substitute in for A_N and rewrite this as

$$\beta(1-\pi_c)^2(A_{N-1} - A_{N-2}) < 1.$$

It will be sufficient to show that $A_{N-1} - A_{N-2} < 1$. To do this I will show that $A_I < 1$. We can do this inductively. Recall that the equation for A_I is

$$A_I = \beta(\pi_c A_{I-1} + (1-\pi_c)A_{I-2}).$$

If $A_{I-1} < 1$ and $A_{I-2} < 1$ then it must be that $A_I < 1$. We only need to prove that A_2 and A_3 (the starting values of A) are less than 1. Note that $A_2 = \pi_c\beta < 1$. We can compute $A_3 = \beta(\pi_c^2\beta + (1-\pi_c))$. To show $A_3 < 1$ notice that if we think of A_3 as a function of π_c it is maximized when $\pi_c = 1/(2\beta)$. If we plug this into A_3 then we find $A_3 = \beta - 1/4 < 1$. This proves the numerator of $V(0)$ is negative. Now let's try to prove the denominator of $V(0)$ is positive.

We can use induction type arguments to do this too. It will be convienient to rewrite the denominator as

$$(1 - B_N)(1 - (1-\pi_c)\beta A_{N-1}) - A_N(\pi_c\beta + (1-\pi_c)\beta B_{N-1})$$

and to prove that $1 - B_N > \pi_c\beta + (1-\pi_c)\beta B_{N-1}$ and $1 - (1-\pi_c)\beta A_{N-1} > A_N$. Let's start with the A inequality. This is equivalent to $\beta A_{N-1} + \beta(1-\pi_c)A_{N-2} < 1$. Let's assume that $\beta A_{I-1} + \beta(1-\pi_c)A_{I-2} < 1$. We want to show this implies that $\beta A_I + \beta(1-\pi_c)A_{I-1} < 1$. Note we can write

$$
\begin{aligned}
\beta A_I + \beta(1-\pi_c)A_{I-1} &= \beta^2\pi_c A_{I-1} + \beta^2(1-\pi_c)A_{I-2} + \beta(1-\pi_c)A_{I-1} \\
&= \beta(1+\pi_c(\beta-1))A_{I-1} + \beta^2(1-\pi_c)A_{N-2} \\
&< \beta A_I + \beta(1-\pi_c)A_{I-1} < 1
\end{aligned}
$$

Then we just have to prove the base case which is that $\beta^2(\pi_c^2\beta + (1-\pi_c)) + \beta^2(1-\pi_c)\pi_c < 1$. This is true since we can reduce the inequality to $\beta^2(1+\pi_c^2(\beta-1)) < 1$.

Now let's do the B inequality. We will rewrite the inequality as $\pi_c\beta + (1-\pi_c)\beta B_{N-1} + B_N < 1$. Assume that $\pi_c\beta + (1-\pi_c)\beta B_{I-2} + B_{I-1} < 1$. Then we can write

$$
\begin{aligned}
B_I + \beta(1-\pi_c)\beta B_{I-1} + \pi_c\beta &= \beta\pi_c B_{I-1} + \beta(1-\pi_c)B_{I-2} + (1-\pi_c)\beta B_{I-1} + \beta\pi_c \\
&= \beta B_{I-1} + \beta(1-\pi_c)B_{I-2} + \beta\pi_c \\
&< \pi_c\beta + (1-\pi_c)\beta B_{I-2} + B_{I-1} < 1
\end{aligned}
$$

Again we have to prove the initial case, which is $\pi_c\beta + (1-\pi_c)^2\beta^2 + \pi_c(1-\pi_c)\beta^2 < 1$. We can reduce this inequality to $\beta(\pi_c + \beta(1-\pi_c)) < 1$, and it is easy to see this will be true.

Next we want to show $V(1) > V(0)$. Since we have proven that the denominators of these terms are positive we need to show that

$$
(1-\pi_c)(B_N - \beta B_{N-1}) + \pi_c(1-\beta) - 1 > (1-\pi_c)(\beta A_{N-1} - A_N) - 1
$$

We can rewrite this inequality as

$$
(1-\pi_c)(\beta(A_{N-1} + B_{N-1}) - (A_N + B_N)) < \pi_c(1-\beta)
$$

It is sufficient to show that $\beta(A_{N-1} + B_{N-1}) - (A_N + B_N) < 0$. Inductive arguments can be used here. First consider the base cases. The base case at 2 boils down to showing $\beta^2 < (1-\pi_c)\beta + \pi_c\beta^2$. This inequality is equivalent to $\beta < (1-\pi_c) + \pi_c\beta$ or $\beta < 1$ (we'd actually get equality is $\pi_c = 1$ but I don't think that is a problem).

Next, we also need to show that $\beta(A_3 + B_3) < A_4 + B_4$. This is straightforward. We know $A_4 + B_4 = \pi_c^2\beta^2 + (1-\pi_c^2)\beta^2$. We need to show that $(1-\pi_c)\beta^2 + \pi_c\beta^3 < \pi_c^2\beta^2 + (1-\pi_c^2)\beta^2$. This inequality reduces to $(\beta-1)\pi_c < (\beta-1)\pi_c^2$. It is easy to see this is true since π_c is a fraction and $\beta < 1$.

To finish we use complete induction. Suppose for all $n < I$ it is the case that $\beta(A_{n-1}+B_{n-1}) < A_n + B_n$. Then we can show

$$
\begin{aligned}
\beta(A_{I-1}+B_{I-1}) &< A_I + B_I \\
\iff \beta(\beta(\pi_c(A_{I-2}+B_{I-2}) + (1-\pi_c)(A_{I-3}+B_{I-3})) &< \beta(\pi_c(A_{I-1}+B_{I-1}) + (1-\pi_c)(A_{I-2}+B_{I-2})) \\
\iff \beta(\pi_c(A_{I-2}+B_{I-2}) + (1-\pi_c)(A_{I-3}+B_{I-3}) &< \pi_c(A_{I-1}+B_{I-1}) + (1-\pi_c)(A_{I-2}+B_{I-2})
\end{aligned}
$$

Our induction assumption implies $\beta(\pi_c(A_{I-2}+B_{I-2}) < \pi_c(A_{I-1}+B_{I-1})$ and $\beta((1-\pi_c)(A_{I-3}+B_{I-3}) < (1-\pi_c)(A_{I-2}+B_{I-2})$.

We also need to show $V(2) > V(1)$. Here is how to do it. We need to prove that $\beta\pi_c V(1) + \beta(1-\pi_c)V(0) > V(1)$. Using our formulas above and the fact that the denominator of the V's is negative we can reduce the inequality to

$$
\begin{aligned}
& \beta^2(1-\pi-c)^2 A_{N-1} + (1-\beta\pi_c)(1-\pi_c)\beta B_{N-1} \\
& > \beta(1-\pi_c)^2 A_N + (1-\beta\pi_c)(1-\pi_c)B_N + (1-\beta\pi_c)(\pi_c(1-\beta)-1) + \beta(1-\pi_c) \\
\iff \quad & (1-\pi_c)([\beta(1-\pi_c)(\beta A_{N-1} - A_N)] + (1-\beta\pi_c)[\beta B_{N-1} - B_N]) > (\beta-1)(\pi_c^2\beta - \pi_c + 1)
\end{aligned}
$$

The right side of the inequality is negative so to make things a bit easier I'm going to multiply both sides by -1 and work with an upper bound. Suppose that

$$
(1-\pi_c)([\beta(1-\pi_c)(A_I - \beta A_{I-1})] + (1-\beta\pi_c)[B_I - \beta B_{I-1}]) < (1-\beta)(\pi_c^2\beta - \pi_c + 1)
$$

for all $I \leq N$. Then we can write the $N+1$ case as

$$
\begin{aligned}
& (1-\pi_c)([\beta(1-\pi_c)(A_{N+1}-\beta A_N)] + (1-\beta\pi_c)[B_{N+1}-\beta B_N]) \\
= \; & (1-\pi_c)([\beta(1-\pi_c)(\beta\pi_c A_N + \beta(1-\pi_c)A_{N-1} - \beta(\beta\pi_c A_{N-1} + \beta(1-\pi_c)A_{N-2})] + \\
& (1-\beta\pi_c)[\beta(1-\pi_c)(\beta\pi_c B_N + \beta(1-\pi_c)B_{N-1} - \beta(\beta\pi_c B_{N-1} + \beta(1-\pi_c)B_{N-2})]) \\
= \; & \beta\pi_c[\beta(1-\pi_c)(A_N-\beta A_{N-1}) + (1-\beta\pi_c)(B_N-\beta B_{N-1})] \\
& \beta(1-\pi_c)[\beta(1-\pi_c)(A_{N-1}-\beta A_{N-2}) + (1-\beta\pi_c)(B_{N-1}-\beta B_{N-2})] \\
< \; & \beta\pi_c(1-\beta)(\pi_c^2\beta-\pi_c+1) + \beta(1-\pi_c)(1-\beta)(\pi_c^2\beta-\pi_c+1) \\
< \; & (1-\beta)(\pi_c^2\beta-\pi_c+1)
\end{aligned}
$$

Last we need to show the base case, for 2, 3 and 4. Recall that $A_2 = \pi_c\beta$, $B_2 = (1-\pi_c)\beta$, $A_3 = \pi_c^2\beta^2 + (1-\pi_c)\beta$, $B_3 = \pi_c(1-\pi_c)\beta^2$, $A_4 = \pi_c^3\beta^3 + 2\pi_c(1-\pi_c)\beta^2$, and $B_4 = \pi_c^2(1-\pi_c)\beta^3 + (1-\pi_c)^2\beta^2$. First we prove the case for time periods 2 and 3. It turns out that the left side of the inequality is

$$
(1-\pi_c)[\beta(1-\pi_c)(\pi_c^2\beta^2 - \pi_c\beta^2 + (1-\pi_c)\beta) + +(1-\beta\pi_c)(1-\pi_c)\beta^2(\pi_c-1)] = 0
$$

. So that part follows. Then we do the same for periods 3 and 4. To start notice that we can write the left hand side of the inequality as

$$
\beta^2(1-\pi_c)^2[\beta(1-\pi_c)(2\pi_c - 1 - (1-\pi_c)\beta) + (1-\beta\pi_c)(\pi_c^2\beta + 1 - \pi_c - \pi_c\beta)]
$$

We want to simplify the stuff inside the square brackets. If you do a bunch of algebra you can reduce this to

$$
\begin{aligned}
& \pi_c\beta - \beta + 2\pi_c\beta^2 - \beta^2 + 1 - \pi_c - \pi_c^3\beta^2 \\
= \; & -(1-\pi_c)\beta + (-\pi_c^3 + 2\pi_c - 1)\beta^2 + (1-\pi_c)
\end{aligned}
$$

We can factor the term multiplying β^2 into $1-\pi_c$ and $\pi_c^2 + \pi_c - 1$. So then then the inequality becomes

$$
\beta^2(1-\pi_c)^3(1-\beta+(\pi_c^2+\pi_c-1)\beta^2) < (1-\beta)(\pi_c^2\beta-\pi_c+1)
$$

Note that we can show that $\pi_c^2+\pi_c-1 < 0$. This is a quadratic that is -1 when π_c is 0, is 0 when π_c is 1, and is strictly increasing in that interval. So it is sufficient to show that $\beta^2(1-\pi_c)^3 < \pi_c^2\beta-\pi_c+1$.

This is straightforward because $\beta^2(1-\pi_c)^3 < 1-\pi_c < \pi_c^2\beta - \pi_c + 1$. The right inequality holds since $\pi_c^2\beta > 0$ and the left one holds since $\beta(1-\pi_c)^2 < 1$.

To show the value function is increasing in inventory for all I we can use the fact we showed $V(1) > V(0)$ and apply complete induction. We assume that our statement is true for all inventory levels n less than I. In other words, if $n < I$ we assume that $V(n) > V(n-1)$ (in particular it means that $V(I-1) > V(I-2) > V(I-3) > \ldots$). We want to show that this implies that $V(I) > V(I-1)$.

Let's start by noticing that $V(I-1) > V(I-2)$ implies the following:

$$V(I-1) - V(I-2) = (A_{I-1} - A_{I-2})V(1) + (B_{I-1} - B_{I-2})V(0) > 0$$

There are two things we will need to complete the proof. First, notice that

$$\begin{aligned} A_I - A_{I-1} &= \beta\pi_c A_{I-1} + \beta(1-\pi_c)A_{I-2} - A_{I-1} \\ &= \beta\pi_c(A_{I-1} - A_{I-2}) + \beta(A_{I-2} - A_{I-1}) \end{aligned}$$

and the same holds true for the B series. The second thing we want to show is that $\beta V(I-2) \geq V(I-1)$. Note that

$$V(I-1) = \beta\pi_c V(I-2) + \beta(1-\pi_c)V(I-3)$$

So we want to argue that

$$V(I-2) \geq \pi_c V(I-2) + (1-\pi_c)V(I-3).$$

We assumed that $V(I-2) > V(I-3)$, so $\pi_c V(I-2) + (1-\pi_c)V(I-3)$ is maximized when $\pi_c = 1$. Then equality holds. Otherwise the inequality must be strict since we are increasing weight on $V(I-3)$.

Now notice that we can write $V(I) - V(I-1)$ as

$$\begin{aligned} V(I) - V(I-1) &= (A_I - A_{I-1})V(1) + (B_I - B_{I-1})V(0) \\ &= (\beta\pi_c(A_{I-1} - A_{I-2}) + \beta(A_{I-2} - A_{I-1}))V(1) + (\beta\pi_c(B_{I-1} - B_{I-2}) + \beta(B_{I-2} - B_{I-1}))V(0) \\ &= \beta\pi_c(V(I-1) - V(I-2)) + \beta V(I-2) - V(I-1) \end{aligned}$$

Our induction assumption was $V(I-1) > V(I-2)$ so $\beta\pi_c(V(I-1)-V(I-2)) > 0$. Additionally we proved above that $\beta V(I-2) - V(I-1) > 0$. Hence, $V(I) - V(I-1) > 0$ and the lemma is proved.

Last we note that all the value functions are continuous in ω_1. This fact will imply that the set of inequalities we proved above will still hold for small values of ω_1.

■

Note that in the Lemma above the bound of $b+3$ is necessary since the consumption shocks can be at most 2. We can use Lemma 4 to prove the Lemma 5, an analog to Lemma 2 from the paper, which is that the proposed policy is optimal.

Lemma 5 *For all $\beta > 0$, it is optimal to purchase only when $I - c < 0$.*

Proof.

We now prove that the proposed policy is optimal. To do this we will show that the value functions we derived in the last lemma are consistent with the optimality conditions. This implies policies derived from them are optimal.

We can prove this lemma inductively. We first note that the consumer will always make a purchase if she is going to run out. To see this, note that the value of purchasing will be $-p-\omega_1+\beta V(b-c_{it})$, while the value from running out and not purchasing is $-\nu+\beta V(0)$. Since the value function is increasing in inventory due to Lemma 1, and we have assumed $p < \nu - \omega_1$, the payoff from purchasing is higher than the payoff from running out. Before we continue we note that this part of the proof is the only place where we use Lemma 1. Lemma 1 is a bit stronger than we need - we really only need it to be the case that $V(b-2) - V(0) \geq 0$ and $V(b-1) - V(0) \geq 0$. These two conditions are harder to prove, but we think they should be true under weaker conditions than what is required to prove Lemma 1.

Next we want to show that the payoff from not purchasing is higher than the payoff from purchasing when the consumer will not run out. It is sufficient to demonstrate that

$$\beta V(0) > -p - \omega_1 + \beta V(b),$$

for I = 0 and

$$\beta V(I) - \omega_1 > -p - \omega_2 + \beta V(b+I),$$

for $I \geq 0$. We begin by showing that $\beta V(0) > -p + \beta V(b)$ and $\beta V(1) - \omega_1 > -p - \omega_2 + \beta V(b+1)$, and then use induction after. To show the first inequality note that following the optimal policy at 0 inventory it must be the case that $V(0) = -p + \beta\pi_c V(b-1) + \beta(1-\pi_c)V(b-2) - \omega_1 = -p + V(b)$.

Thus $V(b) - V(0) = p$. Since $\beta < 1$ and $p + \omega_1$ is positive, it must be that $\beta(V(b) - V(0)) < p$. It is easy to see this implies $\beta V(0) > -p - \omega_1 + \beta V(b)$.

Similar logic can be used to show $\beta V(1) - \omega_1 > -p - \omega_2 + \beta V(b+1)$. First note that we can write $V(1)$ as follows:

$$\begin{aligned} V(1) &= \beta\pi_c V(0) + \beta(1-\pi_c)V(b-1) - (1-\pi_c)p - (1-\pi_c)\omega_1 \\ &= \beta\pi_c(V(b) - p) + \beta(1-\pi_c)V(b-1) - (1-\pi_c)p - (1-\pi_c)\omega_1 \\ &= V(b+1) - (1-\pi_c(1-\beta))p + \pi_c\omega_1 \end{aligned}$$

The inequality is true since $\beta(V(b+1) - V(1)) < (1-\pi_c(1-\beta))p - \pi_c\omega_1 < p + \omega_2 - \omega_1$.

Now we will use induction to prove optimality generally. Suppose that the inequality is true up to $b + I$, in other words that $\beta(V(b+n) - V(n)) < p + \omega_B - \omega_{B-1}$ for $0 \leq n \leq I$. We will work with the value function difference $\beta(V(b+n+1) - V(n+1))$. First, assume no storage cost change between $I+1$, I, and $I-1$. Then we can write this difference as

$$\begin{aligned} \beta(V(b+I+1) - V(I+1)) &= \beta(\pi_c\beta(V(b+I) - V(I)) + (1-\pi_c)\beta(V(b+I-1) - V(I-1)) - (\omega_B - \omega_{B-1})) \\ &< \beta p \\ &< p + \omega_B - \omega_{B-1}, \end{aligned}$$

where the second inequality follows from the induction assumption. Note that our assumption that storage costs are weakly increasing is important here.

Suppose there is a storage cost change between $I+1$ and I. Then the inequalities become

$$\begin{aligned} \beta(V(b+I+1) - V(I+1)) &= \beta(\pi_c\beta(V(b+I) - V(I)) - \pi_c(\omega_B - \omega_{B-1}) + \\ &\quad (1-\pi_c)\beta(V(b+I-1) - V(I-1)) - (1-\pi_c)(\omega_{B-1} - \omega_{B-2})) \\ &< \beta(p + (1-\pi_c)(\omega_B - \omega_{B-1}) + \pi_c(\omega_{B-1} - \omega_{B-2})) \\ &< p + \omega_B - \omega_{B-1}, \end{aligned}$$

The last inequality will follow as a result of weak convexity of the storage cost function. The case where storage costs change between I and $I-1$ is similar. Thus, the policy proposed is optimal.

■

The above two Lemmas lead to an analog of Proposition 1, which we outline below:

Assumptions A1'-A8'

1. Consumption shocks are in the set $\{1,2\}$. The probability that $c_{it}=1$ is π_c.
2. The maximum number of packages that can be purchased in a period is $J=1$.
3. Prices are fixed over time at a level $p>0$.
4. The package size b is greater than or equal to 2.
5. The weight on the error term, η, is small.
6. The stockout cost ν is strictly positive.
7. ω_1 is small, and $p<\nu-\omega_1$.
8. The storage cost function is weakly increasing and weakly convex.

Proposition 3 *If assumptions A1'-A8' hold, and $\eta=0$, then the expected future value of making a purchase, $\beta(V(I+b-c)-V(I-c))$, is decreasing in I for $I\geq c$.*[21] *It is strictly decreasing if $\beta>0$ and 0 if $\beta=0$.*

Proof. The proposition can be proved with Lemma 5 in hand. The inequality we wish to prove is

$$V(I+1+b-c)-V(I+1-c)<V(I+b-c)-V(I-c).$$

Consider the case where $I-c=0$. Using Lemma 2 we can show that

$$\begin{aligned}V(b)-V(0) &= p\\ V(b+1)-V(1) &= (1-\pi_c(1-\beta))p-\pi_c\omega_1\end{aligned}$$

It is clear from the equations above that $V(b+1)-V(1)<V(b)-V(0)$ since $p>0$, $0\leq\beta<1$, and $0\leq\pi_c\leq 1$. We also need to show that $V(b+2)-V(2)<V(b+1)-V(1)$. We can use the fact that $V(2)=\beta\pi_c V(1)+\beta(1-\pi_c)V(0)-\pi_c\omega_1$ to show that $V(b+2)-V(2)=p(\beta\pi_c(1-\pi_c(1-\beta)+\beta(1-\pi_c))-(1+\beta\pi_c^2-\pi_c)\omega_1$. It is sufficient to show the inequality

[21] If the convexity of the storage cost function is violated the proposition still holds, but only for $I\in[0,b-1]$.

$$\beta\pi_c(1-\pi_c(1-\beta)+\beta(1-\pi_c) \quad < \quad 1-\pi_c(1-\beta)$$
$$\iff \beta-\pi_c(1-\beta)(1-\pi_c\beta)<1$$

The second inequality will always be true since $\beta < 1$.

Now we use the induction step. Suppose that $V(n+b)-V(n) < V(n+b)-V(n)$ for all $n=0,...,I+1$ (we can subsume the c into I since it occurs in all the arguments of V). There are a few cases we want to consider. Suppose that there is no storage cost change between $I+2$, I, and $I-1$. In that case the inequality we want to prove is

$$\begin{aligned} & V(I+2+b)-V(I+2)<V(I+1+b)-V(I+1) \\ \iff \quad & \pi_c\beta(V(I+1+b)-V(I+1))+(1-\pi_c)\beta(V(I+b)-V(I)) \\ < \quad & \pi_c\beta(V(I+b)-V(I))+(1-\pi_c)\beta(V(I+b-1)-V(I-1)) \end{aligned}$$

The second inequality in the above equation will be true due to our induction assumption. This proves the proposition up to the level $I+1=b$.

Now suppose we observe a storage cost change between $I+1$ and I, or I and $I-1$. Then, in addition to the above inequality, it will need to be the case that

$$\omega_B-\omega_{B-1}\geq\omega_{B-1}-\omega_{B-2},$$

which is just convexity of storage costs. We note that this is an assumption that is commonly made in prior reasearch on stockpiling. ■

The second proposition we want to prove is that as β rises, the increase in the expected future payoff from making a purchase goes up as inventory goes down for inventory that is below some bound $\overline{I}$. We show this in the following proposition:

Proposition 4 *If assumptions A1-A8 hold, $\eta=0$, then the expected future value of purchase from an increase in inventory, $\beta\left[V(I+b)-V(I)\right)]$, is strictly increasing in β for inventory a for a range $I\in[0,\overline{I}]$, where $\overline{I}\geq 0$.*

Proof. First we note that this payoff will be 0 if $\beta=0$. We first want to show there is some area of the state space where $V(b+I)-V(I)>0$. In this area, it is sufficient to show that $V(b+I)-V(I)$

is weakly increasing in β for the increasingness result to go through. We will use induction to do this.

The base cases are $V(N) - V(0)$ and $V(N+1) - V(0)$. We know that $V(N) - V(0) = p > 0$, so it must be that

$$\frac{\partial V(N) - V(0)}{\partial \beta} = 0.$$

Note that this shows that $\overline{I} \geq 0$. We also know that $V(N+1) - V(1) = (1 - \pi_c(1-\beta))p - (1-\pi_c)\omega_1$. The derivative of this with respect to β is $\pi_c p \geq 0$. As long as $(1 - \pi_c(1-\beta))p - (1-\pi_c)\omega_1 \geq 0$ then the derivative of $\beta(V(b+1) - V(1))$ will be increasing.

Now do the induction step. Suppose that $V(N+n) - V(n)$ is weakly increasing in β for all $n \leq I$, and additionally that $V(N+n) - V(n) > 0$ for all points prior to I. Then let's look at $I+1$. We can write

$$V(N+I+1) - V(I+1) = \beta\pi_c(V_{N+I} - V_I) + \beta(1-\pi_c)(V_{N+I-1} - V_{I-1}) + \Delta\omega,$$

where $\Delta\omega$ is a difference in storage costs (which is not a function of β). Both the terms $V_{N+I} - V_I$ and $V_{N+I-1} - V_{I-1}$ are increasing in β by the induction assumption. This implies that $V(N+I+1) - V(I+1)$ is increasing in β. As long as $V(N+I+1) - V(I+1) > 0$, $\beta(V(N+I+1) - V(I+1))$ will also be increasing in β.

■

Note that the interval over which the expected future payoff is increasing, $[0, \overline{I}]$, will rise storage costs decrease. If there are no storage costs then it can be shown that $\overline{I} = \infty$, because in that case the value function is strictly increasing in inventory, meaning the expected future payoff from purchase is always positive.

12 Appendix: Formal Identification Conditions and Proofs

12.1 Formal Identification Assumptions with Observed Inventory

In this section we provide a set of more formal conditions for identification, to complement the discussion in Section 5. Recall that we wish to identify $M+3$ parameters: the storage costs, the stockout cost, the price coefficient, and the discount factor. If inventory is observed, the consumption rate is identified from the rate at which inventory is depleted from period to period. We maintain Assumptions A1-A7, as well as Assumptions E1 and E2. We add an additional assumption below:

Assumptions I1-I3:

1. The distribution of the error term is type 1 extreme value.

2. Purchase probabilities are observed to the researcher at the following values of inventory:

$$
\begin{aligned}
I_1 &= 0 \\
I_2 &= 1 \\
I_{2+i} &\in [(i-1)b+1, ib], i = 1, ..., M \\
I_{M+3} &\in [2, Mb] \text{ and } I_{M+3} \notin \{I_1, I_2, ..., I_{M+2}\}
\end{aligned}
$$

3. Define $\boldsymbol{I} = (I_1, ..., I_{M+3})$, $\Delta v(I;\boldsymbol{\theta}) = v_1(I;\boldsymbol{\theta}) - v_0(I;\boldsymbol{\theta})$ from equation (15), the vector of $\Delta v(I;\boldsymbol{\theta})$ at $\boldsymbol{I}$ to be $\Delta\boldsymbol{v}(\boldsymbol{I};\boldsymbol{\theta})$, and $\hat{\boldsymbol{P}}_0(\boldsymbol{I};N)$ to be the vector of observed choice probabilities generated at the true parameter vector, $\boldsymbol{\theta}_0$. Then assume that the system of equations defined by

$$\Delta\boldsymbol{v}(\boldsymbol{I};\boldsymbol{\theta}) = \lim_{N\to\infty} \log(\hat{\boldsymbol{P}}_0(\boldsymbol{I};N)) - \log(\boldsymbol{1} - \hat{\boldsymbol{P}}_0(\boldsymbol{I};N)) \tag{21}$$

   has full rank in a neigborhood of $\boldsymbol{\theta}_0$.

We also add the following rank condition:

Proposition 5 *If assumptions A1-A7 and I1-I3 hold then the parameter vector $\boldsymbol{\theta}$ is uniquely identified in a neighborhood of $\boldsymbol{\theta}_0$.*

Proof. The system of equations described in (21) has $M+3$ equations and $M+3$ unknowns, and is full rank. As a result, the Implicit Function Theorem implies that the system can be solved uniquely for $\boldsymbol{\theta}$ in a neighborhood of $\boldsymbol{\theta}_0$, and that solution must be $\boldsymbol{\theta}_0$. In particular, note that under the type 1 extreme value assumption the value functions will be continuous, so that $\lim_{N\to\infty} \log(\hat{\boldsymbol{P}}_0(\boldsymbol{I};N)) - \log(\boldsymbol{1} - \hat{\boldsymbol{P}}_0(\boldsymbol{I};N)) = \Delta\boldsymbol{v}(\boldsymbol{I};\boldsymbol{\theta}_0)$ due to Slutsky's Theorem. Since the system of equations (21) is invertible around $\boldsymbol{\theta}_0$ due to the rank condition, the inverse must be $\boldsymbol{\theta} = \boldsymbol{\theta}_0$. ∎

We note that the assumption of a Type 1 extreme value error can be relaxed - it is sufficient that the error term's CDF is such that one can recover differences in choice specific values from theoretical choice probabilities.

12.2 Formal Identification Conditions with Unobserved Inventory

This section provides some more formal conditions for identification of the parameter vector $\boldsymbol{\theta} = (\alpha, \nu, \beta, \omega_1, ..., \omega_M, \pi_c)$ in the presence of unobserved inventory, as discussed in Section 6. Define the steady state distribution of inventory as $\pi^I(\boldsymbol{\theta})$. Let the indicator $d_t = 1$ if a purchase occurs, and 0 otherwise. Moreover, we will maintain the assumption that prices are fixed over time. Then, define the *purchase hazard*, which is the probability of a purchase occuring in period $t + \tau$ given no purchases in the intervening periods, conditional on observed inventory level I in period t as

$$\phi_\tau(\boldsymbol{\theta}; I) = Prob(d_{t+\tau} = 1 | d_t = 1, d_{t+1} = 0, ..., d_{t+\tau-1} = 0; I_t = I, \boldsymbol{\theta}). \tag{22}$$

Note that to construct the probability in equation (22), we must integrate out over the sequence of consumption shocks that occur in the intervening τ periods between t and $t + \tau$. Since inventory is unobserved, the purchase hazard cannot be computed empirically without a fully specified econometric model. However, the researcher could flexibly estimate the *aggregate purchase hazard*, which is the average probability of a purchase in period $t + \tau$ given a purchase occurs in period t.[22] We define the aggregate purchase hazard as

$$\Phi_\tau(\boldsymbol{\theta}) = \sum_{I=0}^{\bar{I}} \pi^I(\boldsymbol{\theta}) \phi_\tau(\boldsymbol{\theta}; I). \tag{23}$$

We make the following assumptions:

Assumptions U1-U2

1. Assume that the theoretical purchase hazard defined in equation (22) can be computed for $M + 4$ periods, and define this purchase hazard to be

$$\boldsymbol{\Phi}(\boldsymbol{\theta}) = \begin{bmatrix} \Phi_1(\boldsymbol{\theta}) \\ \vdots \\ \Phi_{M+4}(\boldsymbol{\theta}) \end{bmatrix}. \tag{24}$$

   Assume that the the Jacobian of $\boldsymbol{\Phi}(\boldsymbol{\theta})$ is full rank in a neighborhood of the true parameter vector $\boldsymbol{\theta}_0$.

2. The researcher observes the empirical purchase hazard, $\hat{\boldsymbol{\Phi}}(\boldsymbol{\theta}_0; N)$.

Then the following condition can be shown:

[22] Here we assume prices are fixed over time, but price variation is observed in actual data. In this case the research would need to estimate the purchase hazard conditional on the history of observed prices.

Proposition 6 *If Assumptions A1-A7, E1 to E2, and U1-U2 hold then the parameter vector* $\boldsymbol{\theta}$ *can be identified in a neighborhood of* $\boldsymbol{\theta}_0$*.*

Proof. The proof is an application of the Implicit Function Theorem. The empirical purchase hazard, $\boldsymbol{\Phi}(\boldsymbol{\theta}_0; N)$, will be continuous in $\boldsymbol{\theta}_0$ as a result of assumptions E1 and E2, and by Slutsky's Theorem $\lim_{N\to\infty} \hat{\boldsymbol{\Phi}}(\boldsymbol{\theta}_0; N) = \boldsymbol{\Phi}(\boldsymbol{\theta}_0; N)$. The rank condition implies $\boldsymbol{\Phi}(\boldsymbol{\theta})$ is invertible around $\boldsymbol{\theta} = \boldsymbol{\theta}_0$, and so the solution to the system of equations in equation (24) must be $\boldsymbol{\theta} = \boldsymbol{\theta}_0$. ■

The rank condition in U1 may be verified by the researcher. We comment briefly on some additional assumptions which will help the rank condition hold. First, note that in practice the purchase hazard can be computed for some finite number of periods. The larger is the number of periods relative to the number of parameters needed to be estimated, the more overidentifying restrictions will be imposed and as a result the model parameters will be more easily identified. The exclusion restrictions help identification in this sense, as they reduce the number of parameters that need to be estimated.

A second condition that needs to hold is that individuals should not run out of inventory too quickly. To see why a violation of this condition could be problematic, suppose that there are no storage costs, and that $\pi_c = 0$, $b = 2$ and $M = 1$: ie, all individuals use up a package every period and nobody ever holds more than a single package. In this case all inventory will be used at the beginning of the purchase hazard, the purchase hazard will be completely flat, and it will be impossible to identify the discount factor which we argue is identified from the slope of the purchase hazard. Each element of equation (24) will simply be the purchase probability, and the rank of the Jacobian of equation (24) will be 1. Intuitively, in this case we would never know whether the purchase hazard was flat because individuals are myopic, or because the consumption rate is high enough that they do not stockpile.

13 Appendix: Artificial Data Experiment with Continuous Inventory

In this section we describe an additional artificial data experiment where inventory is continuous rather than discrete. We solve for consumer value functions and simulate choices in a market where the utility parameterization is similar to that developed in Section 3. One difference is in how we model storage costs. We run experiments under two different formulations: One is a piecewise linear formulation, where the storage cost is zero for the first three packages and increases linearly by an amount ω for the next packages after that (we assume the maximum number of packages a consumer can purchase is 10. The other assumes that storage costs increase quadratically in the number of packages. We emphasize that in this formulation of the model we continue to maintain the exclusion restriction (except for one instance we discuss at the end of the section): storage costs still depend on the number of packages held, rather than the amount of inventory held within those packages. Prices are assumed to follow a discrete Markov process. We simulate a dataset of 1000 consumers, for 700 periods. We assume in the initial simulation period that all consumers start with inventory of 0. In a data set tracking the behavior of real consumers, consumers will have been making purchases prior to the beginning of the data collection, so initial inventories will be unknown. To mimic this in the simulated data we remove the first 200 periods and estimate the model parameters using the final 500 periods. We estimate the model using the nested fixed point algorithm (Rust 1987) on the simulated data to see how well we can recover the model parameters. Note that when we estimate the model, we need to construct initial inventories in order to evaluate the likelihood. To do this we take the observed 500 periods and split them in half. We assume that in the initial period all consumers have zero inventory, simulate consumption rates, and compute inventory in period 250 as the sum of observed purchases minus consumption rates (where we enforce the restriction that inventory is greater than or equal to zero at each period). We evaluate the likelihood on the final 250 periods for each consumer. To verify that 250 periods is enough to accurately simulate initial inventories, we simulate the model and compute moments of the inventory distribution over time, finding that the inventory distribution looks like it becomes stationary after about 50 periods. We use a simplex algorithm, with a penalty function to enforce nonnegativity constraints, to maximize the likelihood.[23]

We run the artificial data experiment for 8 different parameterizations of the model. The results of the experiment for the first 6 parameterizations are shown in Table 9. The upper half of the table and first 3 columns of the table show, respectively, the estimated parameters, standard errors, and true parameters for the basic specification from the last section. The results here suggest the model parameters are well identified - the estimated parameter values are close to the true values. Turning to the lower half of the table we see under the heading $\beta = 0.6$ that identification is relatively clean for this value of β. The story changes somewhat for the last two sets of estimates, under $\beta = 0.3$ and $\alpha = 0.02$. When $\beta = 0.3$, the estimates of β, the stockout cost ν, and the storage cost ω become significantly worse. The reason for worse identification is similar to what happened in Section 7 when the discount factor was close to zero. In this case consumers behave in a manner that is very close to myopic for low values of β, and so identification of the discount factor becomes more difficult. It is also notable that the storage costs and stockout costs get harder to identify for

[23]We experimented with derivative based methods as well but found the simplex reached slightly higher likelihood values. Occasionally the derivative based methods would stop at the starting points.

lower values of β. The poorer identification of the stockout cost likely occurs because, following our discussion in Section 6.1, when an individual is myopic the stockout cost only affects her purchase decision at the moment she runs out. The poor identification of the storage cost occurs because myopic individuals seldom purchase more than 1 package. It is also notable that when the price coefficient, α, is very low the stockout cost becomes hard to identify, as can be seen in the last 3 columns of the table. The reason this occurs is that price variation helps identify the stockout cost: an individual may decide to pay the stockout cost if prices are sufficiently high. If an individual is price insensitive then this price variation does not affect individuals choices much. Fortunately the discount factor is still well-identified in this situation, suggesting the exclusion restrictions provide identifying power.

Table 10 shows simulated results for a quadratic storage cost, rather than the piecewise linear function. We assume that the linear part of storage costs is 0, and estimate the quadratic coefficient ω. To show that the exclusion restriction has identifying power, we run the artificial data experiment in two scenarios. In the first three columns, we maintain the assumption that storage costs are functions of the number of packages only, enforcing the exclusion restriction. In Column 2, we estimate the parameters in a situation where storage costs are assumed to be continuous rather than discrete (we simulate the model under continuous storage costs and maintain that assumption when we estimate the model). The model specification in Column 2 corresponds to the standard specification used in prior work. All the parameters are still well-identified, suggesting that the nonlinear restrictions from the dynamic model provide some identification. However, it is notable that the standard errors are higher (and sometimes very much higher as in the case of the consumption rates). This suggests that the exclusion restrictions provide identifying variation beyond the functional form restrictions of the model.

14 Appendix: Extension to Inclusive Value Sufficiency

This section describes in detail how the extensions to Inclusive Value Sufficiency (IVS) first proposed in Osborne (2017) can be used to reduce the size of the price state space. Recall that we index brands by k, package sizes by x, and the number of packages chosen by j. Under standard IVS, where error terms are assumed to follow a logit distribution and brand level utilities do not scale with the number of packages chosen, the inclusive value will be a function of both the package size chosen, x, and the number of packages chosen, j:

$$\Omega_{it}(x,j) = \ln\left(\sum_{k}\exp(\xi_{ixk} - \alpha_i p_{ixkt} j)\right).$$

Table 9: Results of Artificial Data Experiment (Quasilinear Storage Cost)

Param	Estimate	S.E.	Truth	Estimate	S.E.	Truth	Estimate	S.E.	Truth
		$\beta = 0.95$			$\beta = 0.99$			β=0.9	
$\underline{c}$	0.001	(0.002)	0	0.001	(0.002)	0	0.001	(0.003)	0
$\overline{c}$	2.012	(0.001)	2	2.011	(0.001)	2	2.014	(0.002)	2
α	0.05	(4.1e-04)	0.05	0.05	(4.4e-04)	0.05	0.05	(4.0e-04)	0.05
ν	0.983	(0.023)	1	0.964	(0.029)	1	0.965	(0.019)	1
β	0.946	(0.002)	0.95	0.985	(0.002)	0.99	0.897	(0.003)	0.9
ω	0.099	(0.003)	0.1	0.101	(0.002)	0.1	0.097	(0.004)	0.1
		β=0.6			β=0.3			α=0.02	
$\underline{c}$	0.004	(0.007)	0	3.9e-04	(0.006)	0	0.001	(0.002)	0
$\overline{c}$	2.004	(0.005)	2	2.002	(0.004)	2	2	(0.001)	2
α	0.05	(3.4e-04)	0.05	0.05	(3.9e-04)	0.05	0.02	(1.4e-04)	0.02
ν	0.967	(0.045)	1	0.876	(0.088)	1	0.565	(0.045)	1
β	0.615	(0.023)	0.6	0.407	(0.064)	0.3	0.944	(0.008)	0.95
ω	0.086	(0.009)	0.1	0.071	(0.013)	0.1	0.11	(0.011)	0.1

Table 10: Estimation with and without the Exclusion Restriction (Quadratic Storage Cost)

Param	Estimate	S.E.	Truth	Estimate	S.E.	Truth
		Discrete			Continuous	
$\underline{c}$	0.001	(0.006)	0	0.001	(0.024)	0
$\overline{c}$	2.013	(0.006)	2	2.013	(0.024)	2
α	0.05	(4.1e-04)	0.05	0.05	(4.3e-04)	0.05
ν	0.975	(0.025)	1	0.987	(0.028)	1
β	0.946	(0.002)	0.95	0.945	(0.002)	0.95
ω	0.005	(1.8e-04)	0.005	0.005	(2.0e-04)	0.005

The number of inclusive values one would have to track will be equal to XK. The idea behind the extension to IVS is to essentially be able to factor j out of the inclusive value. To do this, two assumptions are necessary. First, that brand level utility scales with j, and second, that the error term follows a nested logit distribution where the outer nest corresponds to the choice of package size and number of packages, and the inner nest corresponds to the brand choice. Regarding the first assumption, we assume that the flow utility for choosing j packages of brand k can be written as $\frac{j}{J}\xi_{ixk}$, so that the inclusive value can be written as

$$\Omega_{it}(x,j) = \ln\left(\sum_k \exp\left(\frac{j}{J}[\xi_{ixk} - \alpha_i J p_{ixkt}]\right)\right).$$

Regarding the second assumption, it can be shown using derivations from McFadden (1981) that if a choice-specific error follows an extreme value distribution with scale parameter λ then the expected utility can be written as

$$EU = \lambda \ln\left(\sum_k \exp\left(\frac{j}{J}[\xi_{ixk} - \alpha_i J p_{ixkt}]/\lambda\right)\right).$$

We assume the choice-specific error across package sizes, number of packages and brands follows a nested logit distribution following Cardell (1997). Under this assumption we need the λ parameter to be between 0 and 1, which is why we assume that it is equal to j/J. As a result, the inclusive value can be written as

$$\Omega_{it}(x,j) = \frac{j}{J}\ln\left(\sum_k \exp(\xi_{ixk} - \alpha_i J p_{ixkt})\right).$$

Note that since j can be factored out of the inclusive value, it is not necessary to track a different inclusive value for each j, and it is sufficient to write the inclusive value as $\Omega_{it}(x)$.

15 Appendix: Steps for implementing the IJC algorithm

Before explain the estimation details, we introduce some additional notation. Denote the vector of population-varying parameters drawn in step 1 as $\boldsymbol{\theta}_{i1}$, and the population-fixed parameters in step 2 as $\boldsymbol{\theta}_2$. We assume that the individual- specific parameters are derived from a normal distribution with mean $\boldsymbol{b}'\boldsymbol{Z}_i$ and variance $\boldsymbol{W}$, where $\boldsymbol{Z}_i$ is a matrix of demographic characteristics for household i. Since some of the parameters must be bounded (such as the discount factor or price coefficient) we assume that they are transformations of underlying normal parameters. We assume that the price coefficient, the stockout cost, and the consumption rates are lognormal. The transformation applied to produce the discount factor is $\exp(x)/(1+\exp(x))$, where x is normal. The inventory bound transformation is $M * \exp(x)/(1 + exp(x))$, where the maximum inventory bound M is set to be equivalent to holding 24 of the largest package size of detergent (in terms of volume, this is 4800 liquid ounces of detergent). We will denote the untransformed parameters as $\tilde{\boldsymbol{\theta}}_{i1}$, and the transformed parameters as $\boldsymbol{\theta}_{i1} = T(\tilde{\boldsymbol{\theta}}_{i1})$. Note that we assume that $\tilde{\boldsymbol{\theta}}_{i1} \sim N(\boldsymbol{b}'\boldsymbol{Z}_i, \boldsymbol{W})$.

15.1 Steps 1 to 4: Drawing the model parameters

We use the random walk Metropolis-Hastings Algorithm to implement Step 1 of the Gibbs sampler, and draw the individual specific parameters on a household-by-household basis. To that end we describe how we draw an individual $\boldsymbol{\theta}_{i1}$. Suppose that we are at step g of the Gibbs sampler. First, conditional on the last step's draw of $\tilde{\boldsymbol{\theta}}_{i1}$, which we call $\tilde{\boldsymbol{\theta}}_{i1}^0$, we draw a candidate $\tilde{\boldsymbol{\theta}}_{i1}^1$ from $N(\tilde{\boldsymbol{\theta}}_{i1}^0, \rho_1 \boldsymbol{W}_{g-1})$, where $\boldsymbol{W}_{g-1}$ is last iteration's estimate of the variance matrix. Our new utility parameters will be $\boldsymbol{\theta}_{i1}^1 = T(\tilde{\boldsymbol{\theta}}_{i1}^1)$. We then compute the joint likelihood of brand and size purchase at the old draw and the candidate draw. To implement this we first need an estimate of each consumer's value function. As we describe further in Section 15.2, we compute this estimate by averaging over past value functions, using the nearest neighbor approach of Norets (2009). The choice probability can be written as the probability of the observed brand choice (k_{it}) given package size choice (x_{it}) and number of bottles (j_{it}), multiplied by the probability of the observed size choice. For a given individual the probability of a particular brand choice given their choice of size is

$$Pr(k_{it}|x_{it}, p_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) = \frac{\exp\left(\xi_{x_{it},k_{it}} - J\alpha p_{i,x_{it},k_{it},t}\right)}{\sum_{l \in C_2(x_{it})} \exp\left(\xi_{x_{it},l} - J\alpha p_{i,x_{it},l,t}\right)}.^{24} \tag{25}$$

The probability of a particular size choice can be written independently from the brand choice as

$$Pr(x_{it}, j_{it}|\boldsymbol{\Omega}_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) = \frac{\exp\left(\frac{j_{it}}{J}\Omega_{it}(x_{it}) + \tilde{u}(I_{it}, j_{it}, x_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) + \beta \hat{EV}_i(I_{i,t+1}, \boldsymbol{\Omega}_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2))\right)}{\sum_{(j,x)\in C} \exp\left(\frac{j_{it}}{J}\Omega_{it}(x, j) + \tilde{u}(I_{it}, j, x; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) + \beta \hat{EV}_i(I_{i,t+1}, \boldsymbol{\Omega}_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2)\right)}, \tag{26}$$

where $\hat{EV}(I_{i,t+1}, \boldsymbol{\Omega}_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2)$ is the estimated expected value function, and

$$\tilde{u}(I_{it}, j_{it}, x_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) = -\nu_i \frac{c_i - (I_{it} + x_{it} j_{it})}{c_i} \mathbf{1}\{I_{it} < c_i\}.^{25}$$

[24] Note that the number of bottles purchased, j_{it}, drops out of this choice probability due to the distributional assumptions made on the error term.

[25] The stockout cost enters utility slightly differently in the empirical model, to capture the idea that if a small amount of detergent is left in the bottle the consumer may use a bit of it to do laundry.

Note that to compute this probability we need to compute the inclusive values $\boldsymbol{\Omega}_{it}$, which themselves are functions of $\boldsymbol{\theta}_{i1}$ and $\boldsymbol{\theta}_2$ parameter draws. To construct the estimated value function we will also need to compute the transition process for the inclusive values. We discuss how the inclusive values and their transition process are computed in Section 15.3.

The likelihood used for the Metropolis-Hastings accept-reject step will be

$$L_i(\boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) = \prod_{t=1}^{T_i} Pr(k_{it}|x_{it}, j_{it}, p_{it}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) Pr(x_{it}, j_{it}|\boldsymbol{\Omega}_{it}(1); \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2).$$

The candidate draw will be accepted with probability

$$\frac{L(\boldsymbol{\theta}_{i1}^1, \boldsymbol{\theta}_2)}{L(\boldsymbol{\theta}_{i1}^0, \boldsymbol{\theta}_2)} \frac{k(\tilde{\boldsymbol{\theta}}_{i1}^0)}{k(\tilde{\boldsymbol{\theta}}_{i1}^1)},$$

where k denotes the prior density on $\boldsymbol{\theta}_{i1}$. Under our assumption of normality of the parameters this prior is simply the multivariate normal with mean $\boldsymbol{b}_{g-1}$ and variance $\boldsymbol{W}_{g-1}$.

After drawing the population-varying parameters we draw the mean ($\boldsymbol{b}$) and variance ($\boldsymbol{W}$) parameters that generate them (Steps 2 and 3). Conditional on the $\boldsymbol{\theta}_{i1}$ draws and the demographics $\boldsymbol{Z}_i$, the $\boldsymbol{b}$ parameters are drawn using standard Bayesian regression. We put a relatively diffuse prior on the $\boldsymbol{b}$, using a normal distribution with mean zero and variance matrix of $1000\boldsymbol{I}$.[26]. We also assume a relatively diffuse prior on $\boldsymbol{W}$. If the dimensionality of $\overline{\boldsymbol{\theta}}_{i1,g}$ is K, then the prior variance matrix is set to $\boldsymbol{I} * (K+3) * 0.01$.[27] Given this prior, a posterior draw on the variance matrix can be computed given $\boldsymbol{b}$, $\overline{\boldsymbol{\theta}}_{i1,g}$ and $\boldsymbol{Z}_i$ from an inverse Wishart distribution.[28]

The fourth step is to draw the population-fixed parameters $\boldsymbol{\theta}_2$. This step proceeds in largely the same way as the first step. A candidate draw $\boldsymbol{\theta}_2^1$ is taken from $N(\boldsymbol{\theta}_2^0, \rho_2 \boldsymbol{W}_2)$, and is accepted with probability

$$\frac{\prod_{i=1}^{I} L_i(\boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2^1)}{\prod_{i=1}^{I} L_i(\boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2^0)} \frac{k(\boldsymbol{\theta}_2^0)}{k(\boldsymbol{\theta}_2^1)}.$$

[26]The dimension of $\boldsymbol{I}$ corresponds to $vec(\boldsymbol{b})$

[27]Our choice of prior variance matrix was informed by artificial data experiments, where we generated data from the myopic inventory model with heterogeneous coefficients, and recovered the parameter distributions. We found that when prior variances were too wide the estimator had some difficulty recovering small variance parameters, but scaling down the prior by a factor of about 0.01 mitigated this problem, while still allowing us to recover larger variance parameters.

[28]For more details on the process using to generate the hyperparameters we refer the readers to Rossi, Allenby, and McCulloch (2005). Our code for drawing these parameters is heavily based on the C++ code provided with the book.

We set the prior on $\boldsymbol{\theta}_2$ to be noninformative.

15.2 Step 5: Updating the value function

After a new vector of parameters are drawn, the value functions are updated at the current parameter draw. There are two steps necessary in updating the value function. First, we construct an estimate of the value function at the current parameter draw. Second, we perform a single update to the value function.

We first describe how the estimated value function is constructed. The estimated value function is constructed by integrating over the transition density of inclusive values, and by averaging over past value functions. Importantly, when we average over past value functions, we put more weight on value functions which were computed at parameter draws close to the current draw. For each individual in the data, we store the value function on a grid of 100 inventory points and 100 random price draws (meaning we update the value function for each consumer on a grid of 10,000 points). Index the inventory grid points using s_1 and the price grid points using s_2. The random prices are drawn from the empirical distribution of prices, which we denote as $h(\cdot)$. Since we assume that the value functions are functions of inclusive values, rather than prices, the first thing we do is to compute the inclusive value at each price grid draw which arises at the current parameter draw. We denote this inclusive value as $\boldsymbol{\Omega}_{s_2}$.[29] Once the inclusive values are computed, we compute a set of importance weights, $w_{n_s}(\boldsymbol{\Omega}_{s_2})$, which are used when integrating out the transition probabilities in the value function:

$$w_{n_s}(\boldsymbol{\Omega}_{s_2}) = \frac{F(\boldsymbol{\Omega}_{n_s}|\boldsymbol{\Omega}_{s_2};\boldsymbol{\theta}_{i1},\boldsymbol{\theta}_2)}{h(\boldsymbol{\Omega}_{n_s})},$$

where $F(\boldsymbol{\Omega}_t|\boldsymbol{\Omega}_{t-1};\boldsymbol{\theta}_{i1},\boldsymbol{\theta}_2)$ is the transition density of the inclusive values at the current parameter draw.

The second element is to average over past value functions. For the past $g = 1, ..., G$ Gibbs draws we have $\hat{V}_i$ at each of the inventory states s_1 and the past importance draws on the inclusive values s_2. Denote the saved value functions and parameter draws respectively as $\hat{V}_i^g, \boldsymbol{\theta}_{i1}^g, \boldsymbol{\theta}_2^g, \quad g = 1, ..., G$. We compute kernel weights $\phi_g(\boldsymbol{\theta}_{i1},\boldsymbol{\theta}_2) = \phi([\boldsymbol{\theta}_{i1},\boldsymbol{\theta}_2] - [\boldsymbol{\theta}_{i1}^g,\boldsymbol{\theta}_2^g])$ where $\phi(\cdot)$ is a multivariate normal kernel function. Our value function estimate is then

[29] An alternative approach we experimented with was to propose an importance distribution for inclusive values and to draw inclusive values from that. We found that this approach would sometimes lead to numerical errors when computing transition probabilities, if the current parameter draw was very far from all of the drawn inclusive values. (maybe move this down)

$$\tilde{V}_i(I_{s_1}, \mathbf{\Omega}_{s_2}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) = \sum_{g=1}^{G} \sum_{n_s=1}^{N_s} \frac{\hat{V}^g(I_{s_1}, \mathbf{\Omega}_{n_s}; \boldsymbol{\theta}_{i1}^g, \boldsymbol{\theta}_2^g)\phi_g(\boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) w_{n_s}(\mathbf{\Omega}_{s_2})}{\sum_{g=1}^{G}\sum_{n_s=1}^{N_s} \phi_g(\boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) w_{n_s}(\mathbf{\Omega}_{s_2})}.$$

When we update the value function, we may need to compute the expected value function at inventory points that are not on the grid. To do this we interpolate the $\tilde{V}_i(I_{s_1}, \mathbf{\Omega}_{s_2}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2)$ over inventory states using linear interpolation. Then the expected value function estimate at any inventory point I is

$$\hat{EV}_i(I, \mathbf{\Omega}_{s_2}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) = \tilde{V}_i(\underline{I}'_{s_1}, \mathbf{\Omega}_{s_2}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2) + \frac{I - \underline{I}'_{s_1}}{\overline{I}'_{s_1} - \underline{I}'_{s_1}} \tilde{V}_i(\overline{I}'_{s_1}, \mathbf{\Omega}_{s_2}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2),$$

where $\underline{I}'_{s_1}$ is the largest inventory grid point that is smaller than I and $\overline{I}'_{s_1}$ is the smallest grid point that is larger than I. We will use $\hat{EV}_i(I, \mathbf{\Omega}_{s_2}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2)$ when updating the value function, as we describe in the next paragraph. Before moving on, we also note that when we compute choice probabilities in Section 15.1 we also need to compute the expected value function; for this we use a similar procedure to the above (the main difference is that the value function approximation is evaluated at an inclusive value derived from an observed price, rather than a particular inclusive value grid point).

Denote this value function estimate as $\hat{EV}_i(I, \mathbf{\Omega}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2)$. We will index inventory grid points (which are fixed across Gibbs iterations) with s_1 and inclusive value grid points (which are random) with s_2. Then at a particular grid point s_1, s_2, with state variables $I_{s_1}, \mathbf{\Omega}_{s_2}$, the updated value function is

$$\hat{V}_i(I_{s_1}, \mathbf{\Omega}_{s_2}; \boldsymbol{\theta}_{i1}; \boldsymbol{\theta}_2) = \sum_{(j,x)\in C} \log\left(\exp\left(\hat{\Omega}_{s_2}(x,j) + \tilde{u}(I_{s_1}, j, x; \boldsymbol{\theta}_{i1}; \boldsymbol{\theta}_2) + \beta \hat{EV}_i(I'_{s_1}, \mathbf{\Omega}_{s_1}; \boldsymbol{\theta}_{i1}, \boldsymbol{\theta}_2)\right)\right). \tag{27}$$

15.3 Inclusive value transition process

When we take a new draw on the parameters $\boldsymbol{\theta}_{i1}$ and $\boldsymbol{\theta}_2$ we need to compute new inclusive values, as well as to estimate their transition processes, and the functions for approximating the inclusive values for $j > 1$. Our approach follows Hendel and Nevo (2006a), except we estimate the inclusive value transition process at the individual rather than population level. To do this, first we compute the inclusive values $\Omega_{it}(x)$ for each household and time period. Then, for each household in the data, we run a regression of $\Omega_{it}(x)$ on $\mathbf{\Omega}_{it}$ for each value of x, i.e.:

$$\Omega_{it}(x) = \kappa_{0,i,x} + \sum_{k=1}^{X} \kappa_{k,i,x}\Omega_{i,t-1}(k) + \epsilon_{it,x}.$$

Each regression is run separately for each individual i. We assume that the errors $\boldsymbol{\epsilon}_{it} = (\epsilon_{it,1}, ..., \epsilon_{it,X})$ are normally distributed and i.i.d. over time and individuals. We put no restrictions on the variance matrix of the errors, and we allow the variance to be individual-specific.[30]

15.4 Setup of the Gibbs Sampler

In this section we describe some details of the setup of the Gibbs sampler. The computer code we use is written in R and C++ and designed to take advantage of parallel processing in the value function averaging and updating. Our code for Bayesian estimation makes use of routines from Rossi, Allenby, and McCulloch (2005) for summarizing the model output as well as for drawing the heirarchical parameters. For the Metropolis-Hastings steps in steps 1 and 4 of the Gibbs sampler we need to set the parameters ρ_1 and ρ_2, which control the variance of the random walk process for the population-varying and population-fixed parameters, respectively. Each of these parameters are tuned so that the acceptance rate over the course of the sampler is about 30%. We tune the parameter ρ_1 every iteration: if the fraction of household level parameters that are accepted is above 30%, we increase ρ_1 by 10%; otherwise we decrease it by 10%. For ρ_2, we adjust the parameter every 25 iterations: if the number of acceptances for the past 25 iterations is above 30%, then the ρ_2 parameter is decreased by 25%; otherwise it is increasd by 25%. The ρ parameters move some initially but settle down after about 500 iterations.

To compute the $\boldsymbol{W}_2$ matrix, we estimate the dynamic stockpiling model assuming no unobserved heterogeneity using the 3 step method of Hendel and Nevo (2006a), and compute the inverse information matrix of the likelihood function. We set $\boldsymbol{W}_2$ to be the submatrix of the inverse information matrix corresponding to the parameters which are set to be fixed across the population. We found this procedure for setting $\boldsymbol{W}_2$ worked well in artificial data experiments.

For the value function approximation we choose $G = 30$, and we use a diagonal bandwidth matrix with bandwidth parameter set to $(4/(3G))^{0.2}$. We found these values worked well in artificial data experiments where we tested out the sampler. As we discussed above we evaluate the value function on 100 grid points. Grid points for inventory are chosen between 0 and the maximum inventory value of 4800 ounces, with the first 20 points of the points being clustered equally between 0 and 70 ounces (around the size of a smaller package of detergent) and the rest between 700 and

[30] We experimented with more flexible functional forms for the inclusive value transition process, such one where we used a spline basis. We found that the more flexible models sometimes had issues with overfitting or collinearity; if two variables were collinear then the algorithm we use to run the regression would crash, and occasionally a very flexible functional form would come close to perfectly fitting the data, resulting in an error variance matrix that was close to singular.

4800. We choose more points near zero since the incentive to purchase in advance of running out becomes more important for low inventory levels, and we want to make sure we capture that behavior well.[31] We run the Gibbs sampler for 10,000 iterations. The draws appear to converge at about 4,000 iterations, so we drop the first 4,000 draws to reduce burn-in.

[31] Equivalently, most of the nonlinearity in the value function occurs for low inventory levels, so having more interpolation points in that region of the state space ensures our approximation to the value function is good.