The Value of Demand Visibility in Perishable Goods Supply Chains:

A Simulation Study of Smart Refrigerator Information Systems

Working Paper

January 20, 2026

Abstract

We quantify the value of demand visibility from smart refrigerators in grocery retail. Using simulation, we decompose price effects into components from (i) demand visibility, (ii) markdown pricing, and (iii) consumer picking behavior. Visibility value depends on grocer scale: a small grocer serving 100 households sees 20% price reduction; a larger grocer serving 1,000+ households sees approximately 10%, with pure visibility contributing 5%. This value derives from correlated demand shocks (weather, holidays) that historical forecasting cannot predict but consumption monitoring can observe. The effect is robust to prediction error: imperfect consumption-based forecasts perform nearly as well as perfect demand information from auto-ordering commitments.

1 Introduction

Perishable goods present a fundamental inventory management challenge: grocers must commit to orders days before knowing actual demand, yet unsold inventory spoils. This uncertainty manifests as waste, which grocers recover through higher margins. The resulting price inflation is substantial—industry estimates suggest 30–40% of perishable produce is wasted in the supply chain, with costs passed to consumers. Industry initiatives to reduce this waste have documented significant returns: Champions 12.3, a coalition tracking progress toward UN Sustainable Development Goal 12.3, reports a median 14:1 ROI on food waste reduction investments across 1,200 business sites, with leading retailers achieving 25–30% waste reductions (Champions 12.3, 2017).

We study whether *demand visibility*—advance knowledge of household consumption needs—can reduce this inefficiency. The motivating technology is a smart refrigerator that monitors household inventory levels and predicts restocking needs. By aggregating these signals across households, grocers could observe demand before placing supplier orders, reducing both waste and the precautionary margins that fund it.

Our contribution is threefold. First, we construct a simulation model that captures the key economic mechanisms: stochastic household demand, perishable inventory with quality degradation, grocer ordering under uncertainty, and cost-plus pricing that passes waste costs to consumers. Second, we use factorial experimental design to decompose the total value of the smart fridge system into components from demand visibility versus operational improvements (markdown pricing, consumer behavior). Third, we establish that the value of visibility is robust to prediction error, suggesting that "soft" predictions based on consumption monitoring may be as valuable as "hard" consumer commitments.

1.1 Related Work

Our work builds on two foundational literatures: inventory management under demand uncertainty and supply chain information sharing.

1.1.1 Inventory Policy Under Uncertainty

The canonical framework for inventory management under stochastic demand derives from Arrow, Harris, and Marschak (1951), who establish the optimality of (s, S) policies for single-product inventory problems with linear costs. In their formulation, the firm faces a trade-off between holding costs (from excess inventory) and shortage costs (from unmet demand). The optimal policy is characterized by a reorder point s and order-up-to level S: when inventory falls below s, order enough to restore to S.

Our implementation follows this structure. The grocer reviews inventory daily and orders to an order-up-to level that balances service level (avoiding stockouts) against waste from overstocking perishables. The key difference from the classical model is that our "shortage cost" includes lost sales plus consumer substitution to competitors, while our "holding cost" includes actual spoilage—not merely opportunity cost of capital.

1.1.2 The Value of Information Sharing

A central theoretical question is: under what conditions does demand information create value? Lee, So, and Tang (2000) analyze a two-stage supply chain where the retailer shares demand observations with the manufacturer. They formalize an intuitive result: *when demand is i.i.d., information sharing has zero value*—observing current demand provides no signal about future demand. Information sharing is valuable only when demand is autocorrelated. Modeling demand as AR(1) with autocorrelation ρ, they show that value increases monotonically in ρ, achieving approximately 20% inventory reduction at $\rho = 0.7$. This reduction matters because inventory carries costs: working capital tied up in stock, warehouse space, and—critically for perishables—spoilage risk. A 20% inventory reduction translates directly to lower safety stock requirements, freeing capital and reducing waste.

Gallego and Özer (2001) study a distinct information structure: customers who place orders in advance of their needs, giving the firm direct knowledge of future demand rather than observations to improve forecasts. Their work establishes that state-dependent (s, S) policies are optimal when the system state incorporates this advance demand information (ADI), and characterizes conditions under which ADI creates operational value.

1.1.3 How Our Model Differs

The smart fridge provides *direct visibility into future demand*, not improved forecasts from demand observations. This visibility operates through two mechanisms:

1. **Consumption observation**: The fridge monitors inventory levels and consumption rates. Seeing that milk is at 15% remaining and depleting at a known rate, the system predicts when replenishment will be needed. This is still forecasting, but from a richer signal than purchase history: the physical state of the consumption cycle, not merely when the household last transacted.

2. **Full autonomy (hard commits)**: The consumer grants the fridge authority to auto-order. When inventory drops below a threshold, the fridge automatically

schedules replenishment. This converts prediction into *committed demand*: the grocer knows with certainty that household #47 will receive milk on Tuesday. This is oracle knowledge, not forecasting.

As we show in Section 4.5, even noisy consumption-based predictions capture most of the value of hard commits, suggesting that full autonomy is not strictly necessary.

The key implication: *the value of smart fridge visibility does not depend on demand autocorrelation.* Lee, So, and Tang (2000)'s zero-value result for i.i.d. demand applies when information means observing past demand to forecast future demand. The smart fridge observes current inventory state—a signal that purchase history cannot provide—and with full autonomy, bypasses forecasting entirely.

We validate this using the Instacart dataset (3.4 million orders, 206,000 users). For staple essentials, *what* households buy is highly predictable (reorder rates exceed 70%), but *when* they buy exhibits high variability: inter-purchase intervals have mean 10.3 days with standard deviation 8.7 days (CV = 0.84). This timing uncertainty is what inventory visibility solves. The slight negative autocorrelation in repurchase intervals ($\rho \approx -0.17$) reflects inventory depletion cycles—a household that last bought milk 12 days ago is more likely to buy tomorrow—but this signal is directly observable from current inventory state, not inferred from purchase patterns.

2 Model

We model a discrete-time economy with three agent types: households, a grocer, and suppliers. Time is indexed by $t \in \{0, 1, \ldots, T\}$ where each period represents one day.

2.1 Products

Let M denote the number of SKUs, indexed $j \in \{1, \ldots, M\}$. Each SKU j has attributes:

- $\ell_j \in \mathbb{N}$: supplier lead time (days between order and delivery)
- $\tau_j \in \mathbb{N} \cup \{\infty\}$: shelf life (days until spoilage; ∞ for non-perishables)
- $c_j \in \mathbb{R}_+$: wholesale cost per unit

We partition SKUs into categories (produce, dairy, meat, frozen, packaged) with category-specific parameter distributions calibrated to grocery industry data. Approximately 60% of SKUs are perishable ($\tau_j < \infty$).

2.2 Households

Let N denote the number of households, indexed $i \in \{1, \ldots, N\}$. Each household i has a preference mapping $\mathcal{P}_i$ from a subset of SKUs to consumption parameters:

$$\mathcal{P}_i : \mathcal{J}_i \to \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{N}, \quad \mathcal{J}_i \subseteq \{1, \ldots, M\}$$

where $\mathcal{P}_i(j) = (\mu_{ij}, \sigma_{ij}, q_{ij})$ specifies:

- μ_{ij}: mean days between purchases
- σ_{ij}: standard deviation of inter-purchase times
- q_{ij}: typical purchase quantity

We write $j \in \mathcal{J}_i$ to indicate that household i purchases SKU j.

Inter-purchase times follow a Gamma distribution with shape $\alpha = (\mu_{ij}/\sigma_{ij})^2$ and scale $\theta = \sigma_{ij}^2/\mu_{ij}$. We calibrate $\mu = 10$ days and $\sigma = 9$ days based on grocery purchase frequency data, yielding a right-skewed distribution with substantial variance.

2.2.1 Demand Generation

We pre-compute a deterministic demand schedule $D : \{0, \ldots, T\} \times \{1, \ldots, M\} \to \mathbb{N}$ where $D(t, j)$ is aggregate demand for SKU j on day t. This schedule incorporates:

1. **Baseline stochasticity**: Each household's purchase timing is drawn from their Gamma distribution.
2. **Day-of-week effects**: Weekend demand is scaled by 1.3; Monday demand by 0.8.
3. **Correlated shocks**: On 20% of days, a global demand multiplier is drawn uniformly from $[0.7, 1.5]$, representing weather or local events affecting all households.

4. **Holiday effects**: Major holidays (Thanksgiving, Christmas, July 4th) receive demand multipliers of 1.3–1.8.

Pre-computing demand ensures identical realized demand across regimes, isolating the effect of information rather than demand realization.

2.3 Grocer

The grocer maintains inventory, places orders, sets prices, and fulfills household demand. Let $I_j(t)$ denote inventory of SKU j at time t, and $O_j(t)$ denote the order quantity placed on day t (arriving on day $t + \ell_j$).

2.3.1 Inventory Dynamics

Inventory evolves as:

$$I_j(t+1) = I_j(t) + R_j(t) - F_j(t) - W_j(t)$$

where $R_j(t)$ is units received from supplier orders, $F_j(t)$ is units fulfilled (sold), and $W_j(t)$ is units wasted due to spoilage.

Each inventory unit carries a timestamp of receipt. Quality degrades non-linearly:

$$\text{quality}(a) = 1 - \left(\frac{a}{\tau_j}\right)^{1.5}$$

where a is days since receipt. Units with quality below 0.2 are discarded as waste.

2.3.2 Ordering Policy

We implement an (s, S) inventory policy. On each day, the grocer reviews inventory position (on-hand plus in-transit) and orders to restore to target level S_j if position falls below reorder point s_j.

Regime 1 (Status Quo). The grocer forecasts demand using Facebook's Prophet algorithm with weekly seasonality, trained on observed sales history. The order-up-to level is:

$$S_j = \hat{d}_j \cdot (\ell_j + r) + z \cdot \hat{\sigma}_j \cdot \sqrt{\ell_j + r}$$

where $\hat{d}_j$ is forecast daily demand, $r = 1$ is the review period (daily ordering), $\hat{\sigma}_j$ is forecast standard deviation, and $z = 1.65$ corresponds to 95% service level.

A key implementation detail is *censored demand*: when stockouts occur, observed sales understate true demand. We impute censored demand using Prophet's forecast for stockout periods.

Regime 2 (Smart Fridge). The grocer observes future demand $D(t', j)$ for $t' \in \{t+1, \dots, t+k\}$ where k is the commitment horizon. When lead time $\ell_j \leq k$, the grocer has perfect demand visibility and orders:

$$O_j(t) = \max \left\{ 0, \sum_{t'=t+1}^{t+\ell_j+1} D(t', j) + \text{buffer} - I_j(t) - \text{in-transit}_j(t) \right\}$$

with a 5% buffer for timing mismatches. When $\ell_j > k$, the grocer combines known demand for days $t+1$ through $t+k$ with forecasted demand for days $t+k+1$ through $t+\ell_j$.

2.3.3 Pricing

Prices follow a cost-plus model:

$$p_j = c_j \cdot (1 + m_{\text{waste}} + m_{\text{comp}} + m_{\text{profit}})$$

where $m_{\text{waste}} = w/(1-w)$ recovers waste costs (with w being historical waste rate), $m_{\text{comp}} = 0.05$ is competitive buffer, and $m_{\text{profit}} = 0.10$ is target profit margin.

2.3.4 Markdown Pricing

When enabled, items with quality between 0.2 and 0.6 are marked down by 50%. This allows recovery of some value from aging inventory rather than complete spoilage.

2.4 Consumer Behavior

When households shop, they select items from available inventory. We model heterogeneous picking behavior:

- Fraction λ of consumers are "freshness seekers" who select the newest items (LIFO: last in, first out)
- Fraction $1 - \lambda$ select the oldest items (FIFO: first in, first out), either by habit or to claim markdown prices

This mixing is economically important: pure LIFO behavior causes older inventory to age out, increasing waste even when total inventory is adequate.

3 Experimental Design

We employ a $2 \times 2 \times 2$ factorial design crossing three factors:

1. **Information regime**: Prophet forecasting (Regime 1) vs. smart fridge visibility (Regime 2)
2. **Markdown pricing**: Disabled vs. enabled
3. **Consumer picking**: 100% LIFO ($\lambda = 1$) vs. 60% LIFO ($\lambda = 0.6$)

This yields four named conditions:

Condition	Description	Regime	Markdown	LIFO fraction
A	Prophet / No markdown / LIFO	1	Off	1.0
B	Prophet / Markdown / Mixed	1	On	0.6
C	Smart Fridge / No markdown / LIFO	2	Off	1.0
D	Smart Fridge / Markdown / Mixed	2	On	0.6

Condition A represents a "broken" baseline with suboptimal operations. Condition B represents best practices achievable without smart fridge technology. Condition D represents the full smart fridge system.

3.1 Effect Decomposition

The factorial structure enables decomposition of total effects:

- **Pure visibility effect** (A $\rightarrow$ C): Holding markdown and picking constant, what does visibility alone contribute?
- **Operational effect** (A $\rightarrow$ B): Holding information constant, what do markdown pricing and realistic picking contribute?
- **Total effect** (A $\rightarrow$ D): Combined effect of all improvements.
- **Incremental visibility effect** (B $\rightarrow$ D): The policy-relevant comparison—what does visibility add to an already-optimized baseline?

3.2 Simulation Parameters

Each condition is simulated with:

- $N = 100$ households
- $M = 50$ SKUs (60% perishable)
- $T = 180$ days (6 months)
- Commitment horizon $k = 5$ days
- 5 independent Monte Carlo replications (seeds 42–46)
- 30-day warmup period excluded from metrics to ensure steady-state measurement

Initial inventory is calibrated to expected demand based on SKU popularity and household count, avoiding initialization artifacts.

3.3 Outcome Variables

For each simulation run, we record:

- **Waste rate**: Spoiled units / ordered units
- **Fulfillment rate**: Fulfilled units / demanded units
- **Price per unit**: Total consumer spending / units purchased

- **Waste per household-day**: Absolute waste metric, units/(households × days)

Price per unit is our primary outcome, as it directly measures consumer welfare impact.

4 Results

4.1 Main Effects

Table 1 reports mean outcomes across 5 replications, using steady-state metrics (post-warmup).

Table 1: Factorial Experiment Results (5 seeds, steady-state metrics)

Condition	SS Waste Rate	Fulfillment	$/Unit
A: Prophet/NoMD/LIFO	2.9%	87.1%	$8.45
B: Prophet/MD/Mixed	0.0%	87.8%	$7.27
C: SmartFridge/NoMD/LIFO	12.3%	99.6%	$8.72
D: SmartFridge/MD/Mixed	0.1%	100.0%	$6.53

Several patterns emerge. First, the smart fridge regime dramatically improves fulfillment (99%+ vs. 87–88%), confirming that demand visibility reduces stockouts. Second, waste rates in Regime 2 are *higher* than Regime 1 when comparing analogous conditions (C vs. A, D vs. B). We discuss this counterintuitive finding below. Third, despite higher waste rates, Regime 2 achieves lower prices because fulfillment improvements and markdown efficiency dominate.

4.2 Effect Decomposition

Table 2 decomposes price savings relative to condition A.

Two key findings emerge. First, the B → D comparison shows demand visibility contributes 10.2% price savings *beyond* what an optimized baseline achieves—a meaningful incremental effect. Second, and more strikingly, pure visibility without markdown pricing (A → C) actually *increases* prices by 3.2%. This reveals that visibility's value is contingent on complementary markdown pricing to clear aging inventory.

Table 2: Effect Decomposition

Comparison	Price Savings	Interpretation
A $\to$ C	-3.2%	Pure visibility (no markdown)
A $\to$ B	14.0%	Markdown + picking effect
A $\to$ D	22.6%	Total effect
B $\to$ D	10.2%	Incremental visibility effect

4.3 Statistical Inference

For each comparison, we compute per-seed savings and construct confidence intervals. Let $p_X^{(s)}$ and $p_Y^{(s)}$ denote price per unit in conditions X and Y for seed s. The savings percentage is:

$$\delta^{(s)} = \frac{p_X^{(s)} - p_Y^{(s)}}{p_X^{(s)}} \times 100$$

For the incremental visibility effect (B $\to$ D), we obtain mean savings of 10.2% with 95% confidence interval [7.5%, 12.0%]. For the total system effect (A $\to$ D), we obtain mean savings of 22.6% with 95% confidence interval [19.6%, 27.9%].

Notably, pure visibility without markdown (A $\to$ C) shows a *negative* effect, with prices increasing by 3.6% on average (95% CI: -18.7% to $+3.7\%$). The wide confidence interval reflects the interaction between visibility and markdown pricing—visibility alone is not sufficient.

4.4 The Waste Paradox

A striking finding is that steady-state waste *increases* with demand visibility when markdown pricing is disabled: condition C exhibits 12.3% waste versus 2.9% for condition A. This is a genuine phenomenon, not a measurement artifact.

The mechanism is as follows. With demand visibility, the grocer achieves near-perfect fulfillment (99.6%) by maintaining adequate inventory. However, with 100% LIFO consumer picking, the newest items are selected first, causing older inventory to age out. In the status quo, frequent stockouts (87.1% fulfillment) "clear" aging inventory before it spoils. Perfect demand visibility ensures the grocer orders the correct *quantity*, but cannot eliminate the *timing mismatch* between item arrival and consumer pickup—items age on the shelf while awaiting their committed buyer, and freshness-seeking consumers bypass older stock.

This finding has important implications. First, waste rate is a poor metric for evaluating information systems; fulfillment and price are more appropriate. Second, the value of visibility comes not from waste reduction per se, but from enabling higher fulfillment at lower prices. Third, markdown pricing is *essential*—it steers some consumers toward older inventory, reducing waste from 12.3% (C) to 0.1% (D). Without markdown, visibility alone actually increases prices (the negative A $\rightarrow$ C effect).

4.5 Robustness: Hard Commits vs. Predictions

A natural question is whether the value of visibility requires explicit consumer commitments or whether predictions suffice. We compare two information models:

1. **Hard commits**: The grocer observes exact future demand $D(t', j)$.
2. **Noisy predictions**: The grocer observes $\hat{D}(t', j) = D(t', j) \cdot (1 + \epsilon)$ where $\epsilon \sim N(0, 0.15)$.

Table 3 compares B $\rightarrow$ D savings under each model (5 seeds).

Table 3: Hard Commits vs. Noisy Predictions

Information Model	B $\rightarrow$ D Savings	Fulfillment (D)
15% prediction error	9.3%	99.3–99.5%
Hard commits (0% error)	10.2%	99.8–100.0%
Difference	+0.9pp	+0.5pp

Hard commits yield only 0.9 percentage points additional savings. The grocer's safety stock buffers absorb prediction error effectively. This suggests that consumption-based predictions—easier to implement than commitment systems—capture nearly all the value of demand visibility.

4.6 Scale Dependence

A critical question for practical deployment is how visibility value scales with the number of households served. We conducted a scaling analysis across six deployment sizes: 100, 250, 500, 1,000, 2,000, and 5,000 households, with 10 Monte Carlo replications per scale.

Table 4: Visibility Value by Deployment Scale (10 seeds each)

Households	A $\rightarrow$ D Savings	B $\rightarrow$ D Savings	Std. Dev.
100	20.4%	9.2%	$\pm$3.3%
250	12.4%	5.6%	$\pm$1.3%
500	11.0%	5.4%	$\pm$1.5%
1,000	11.1%	5.6%	$\pm$1.3%
2,000	11.4%	5.9%	$\pm$1.4%
5,000	11.5%	6.0%	$\pm$1.5%

The results reveal a striking pattern: visibility value decreases sharply from 100 to 500 households, then stabilizes. We fit an asymptotic model of the form:

$$\text{savings}(n) = a + \frac{b}{\sqrt{n}}$$

where a represents the asymptotic value that persists at any scale, and $b/\sqrt{n}$ captures idiosyncratic variance that averages out. For the pure visibility effect (B $\rightarrow$ D), we estimate $a = 4.7\%$ with $R^2 = 0.58$.

Interpretation. At small scale (100 households), individual household demand is "lumpy"—a single household's party or vacation creates substantial demand variance. The smart fridge observes these idiosyncratic events in real-time, while Prophet forecasts from smoother historical averages. This information asymmetry is valuable.

At large scale, the law of large numbers operates: idiosyncratic household events average out, making aggregate demand smoother and more predictable. Prophet's historical forecasts become reasonably accurate for aggregate demand, reducing the marginal value of visibility.

The asymptotic value of approximately 5% arises from *correlated* demand shocks—weather events, holidays, and other factors that affect all households simultaneously. These shocks do not average out with scale. Prophet cannot predict them from historical data, but the smart fridge observes them in real-time as consumption patterns shift across households.

For practical deployment planning, this implies:

- Total system savings (vs. unoptimized baseline): approximately 10% at scale
- Pure visibility value (vs. optimized baseline): approximately 5% at scale

- Variance in outcomes decreases with scale, making results more predictable

5 Discussion

5.1 Economic Interpretation

Our results support the hypothesis that demand visibility creates consumer surplus in perishable goods markets, with two important caveats: (1) visibility must be coupled with markdown pricing to realize its value, and (2) the magnitude of savings is scale-dependent. The mechanism operates through three complementary channels:

1. **Fulfillment improvement**: Visibility enables 99–100% fulfillment versus 87–88% under forecasting, reducing deadweight loss from stockouts.
2. **Waste reduction through markdown**: With markdown pricing enabled, steady-state waste falls to near zero (0.1%), as aging inventory is sold at discount before spoiling.
3. **Margin compression**: Lower waste rates and higher fulfillment allow lower cost-plus margins.

At realistic deployment scale (1,000+ households), the total system effect ($A \to D$) converges to approximately 10% price reduction, with pure visibility contributing roughly 5%. For a household spending \$200/week on groceries with 40% perishables, annual savings would be approximately \$400 from the full system, or \$200 from visibility alone.

These are meaningful but modest savings—far from transformative. The value proposition of smart fridge technology should therefore emphasize the combination of (i) consistent 5% cost savings, (ii) near-perfect fulfillment eliminating stockout frustration, and (iii) convenience from automated restocking. The savings alone may not justify adoption; the full value proposition is multidimensional.

5.2 Connection to Theory

Our setting differs fundamentally from the classical information sharing framework of Lee, So, and Tang (2000). Their model analyzes information sharing as a *forecasting* *improvement*: observing past demand helps predict future demand when demand is autocorrelated. The value of information in their framework depends entirely on autocorrelation structure—with i.i.d. demand, information sharing has zero value because past observations provide no signal about the future.

The smart fridge operates through an entirely different mechanism: *direct observation of future demand*. By monitoring household inventory states, the grocer gains oracle knowledge of what will be purchased, not better forecasts based on what was purchased. A household with one day of milk remaining will buy milk tomorrow—this is knowable regardless of whether purchase history exhibits autocorrelation. The smart fridge bypasses the forecasting problem entirely, which is why autocorrelation structure is irrelevant to its value.

The mechanism in our simulation is therefore distinct. The incremental price savings from visibility does not arise primarily from better aggregate forecasting, but from two complementary channels: (i) precise timing of individual household needs enables near-perfect fulfillment (100% vs 88%), eliminating deadweight loss from stockouts; and (ii) the combination of visibility and markdown pricing virtually eliminates waste (0.1% vs higher rates without visibility).

However, our scaling analysis (Section 4.6) reveals an important nuance: these benefits are most pronounced at small scale, where individual household events create substantial demand variance. At large scale, idiosyncratic household events average out, and Prophet's historical forecasts become reasonably accurate for aggregate demand. The asymptotic visibility value of approximately 5% arises specifically from *correlated* shocks—weather, holidays, and other factors affecting all households simultaneously—that Prophet cannot predict but real-time monitoring can observe.

Critically, visibility *alone* is not sufficient—without markdown pricing, visibility actually increases prices (the negative A $\rightarrow$ C effect). The value of visibility is realized through the synergy with markdown pricing, which clears aging inventory that LIFO consumer picking would otherwise strand.

A corollary is that our results should *not* be interpreted as universal. The smart fridge value proposition is specific to categories with persistent consumption patterns—staples, perishables, and household essentials—where individual households have regular replenishment needs. Markets with genuinely unpredictable demand (fashion, novelty items) or infrequent purchases would not benefit similarly.

5.3 Limitations

Several limitations warrant discussion.

Single grocer. We model a monopolist grocer. Competition would accelerate pass-through of cost savings to consumers but might reduce incentives to invest in visibility infrastructure.

Homogeneous preferences. Households have idiosyncratic preference sets but homogeneous parameters. Heterogeneity in price sensitivity or freshness preferences could affect results.

No strategic behavior. Households do not respond strategically to the smart fridge system (e.g., by gaming predictions). In practice, such responses might erode information value.

Fixed commitment horizon. We fix $k = 5$ days. The value of visibility likely varies with k and with the distribution of supplier lead times.

Conservative calibration to essentials. Our gamma distribution for inter-purchase intervals is calibrated to essential groceries (milk, eggs, bread), which exhibit the most routine purchasing patterns (CV = 0.84). Non-essential categories—meats, frozen foods, specialty items—typically show higher timing variability due to event-driven consumption (e.g., barbecues, dinner parties, seasonal cooking). Since the value of visibility increases with timing uncertainty, our estimates represent a lower bound. Category-specific modeling would likely show larger visibility effects for non-routine purchases, though essentials dominate grocery volume.

5.4 Implementation Considerations

Our finding that noisy predictions perform nearly as well as hard commits has practical implications. A smart fridge system need not require consumers to make binding commitments—passive consumption monitoring with predictive algorithms may suffice. This substantially simplifies the user experience and reduces friction to adoption.

6 Conclusion

We have quantified the value of demand visibility in perishable goods supply chains, with particular attention to how this value scales with deployment size. Our factorial simulation study finds that smart fridge technology, combined with markdown pricing, yields meaningful but scale-dependent consumer price savings. At small scale (100 households), total system savings reach 20%; at realistic deployment scale (1,000+ households), this converges to approximately 10%. The pure visibility effect—incremental savings from demand information beyond operational improvements—is approximately 5% at scale.

This scaling behavior has a clear interpretation: at small scale, individual household demand is lumpy and unpredictable, making real-time visibility valuable. At large scale, the law of large numbers smooths aggregate demand, reducing the marginal value of observing individual households. The asymptotic 5% value arises specifically from *correlated* demand shocks (weather, holidays) that affect all households simultaneously and cannot be predicted from historical data.

A key finding is that visibility *alone* is not sufficient—without markdown pricing, visibility actually increases prices. The value of visibility is realized through synergy with markdown pricing, which clears aging inventory that consumer freshness-seeking behavior (LIFO picking) would otherwise strand. These effects are robust to 15% prediction error, suggesting that consumption-based predictions capture nearly all the value of explicit commitments.

The key empirical insight underlying these results is the distinction between *what* and *when*. Using Instacart transaction data (3.4 million orders, 206,000 users), we find that *what* households buy is highly predictable—staple essentials have reorder rates exceeding 70%. But *when* they buy exhibits high variability: inter-purchase intervals have a coefficient of variation of 0.84, meaning timing is effectively unpredictable from purchase history alone. This timing uncertainty is exactly what inventory visibility resolves. The smart fridge observes current inventory state, not just past purchases, enabling precise demand timing that history-based forecasting cannot achieve.

For practitioners considering smart fridge technology, our results suggest tempering expectations: the economic value at scale is approximately 5% in pure visibility savings, or 10% total system savings including operational improvements. These are meaningful but not transformative. The business case for smart fridge technology likely rests on the combination of (i) consistent cost savings, (ii) near-perfect fulfillment eliminating stockout frustration, and (iii) convenience from automated restocking—not on dramatic price reductions alone.

A Simulation Algorithm

Algorithm 1 Daily Simulation Step

for each household $i \in \{1, \ldots, N\}$ **do**
    Consume from household refrigerator inventory
end for
Grocer receives supplier deliveries due today
for each household $i \in \{1, \ldots, N\}$ **do**
    Retrieve pre-computed demand $D(t, j)$ for household i
    if demand is non-empty **then**
        Grocer fulfills order (LIFO/FIFO picking)
        Household receives groceries
    end if
end for
Grocer calculates spoilage (quality < 0.2)
Grocer marks down items with quality $\in [0.2, 0.6]$
Grocer finalizes daily demand history
Grocer places orders to supplier
if $t \mod 7 = 0$ **then**
    Grocer updates prices based on waste rates
end if

B Parameter Calibration

References

Arrow, Kenneth J., Theodore Harris, and Jacob Marschak. Optimal inventory policy. *Econometrica*, 19(3):250–272, 1951.

Champions 12.3. The business case for reducing food loss and waste. Technical report, World Resources Institute, March 2017.

Gallego, Guillermo and Özalp Özer. Integrating replenishment decisions with advance demand information. *Management Science*, 47(10):1344–1360, 2001.

Lee, Hau L., Kut C. So, and Christopher S. Tang. The value of information sharing in a two-level supply chain. *Management Science*, 46(5):626–643, 2000.

Table 5: Simulation Parameters

Parameter	Description	Value
N	Number of households	100
M	Number of SKUs	50
T	Simulation days	180
k	Commitment horizon	5 days
μ	Mean inter-purchase days	10
σ	Std inter-purchase days	9
z	Safety stock z-score	1.65
m_{comp}	Competition margin	5%
m_{profit}	Profit margin	10%
λ	LIFO fraction (mixed)	0.6