Decision Making and Analysis
https://ojs.luminescience.cn/DMA

LUMINESCIENCE
https://luminescience.cn

Original Research

Identifying the effective cost ratio and a model-consistent shadow carbon price in deteriorating item inventory systems: a deterministic inverse optimisation approach

Janardan Behera

Department of Statistics, Ravenshaw University, Cuttack (Odisha), 753003, India

Corresponding to: Janardan Behera, Email: janardan@ravenshawuniversity.ac.in

Abstract: Firms in perishable goods supply chains leave a detailed trail of replenishment decisions, yet the cost structure behind those decisions, and the value a firm place on its own carbon emissions, stay hidden. This paper develops a deterministic inverse optimisation framework that reads part of that hidden structure back out of observed behaviour. A forward inventory model for a deteriorating item shows that ordering, holding, spoilage and carbon costs collapse into an effective setup term and an effective holding term, so behaviour identifies only a single cost ratio and the deterioration rate; individual accounting costs are not recoverable from decisions alone. A two-stage estimator recovers the deterioration rate from the inventory balance linking orders to sales and the cost ratio from the optimality condition, then backs out a model-consistent shadow carbon price from the ratio and the firm's known accounting costs. We prove which quantities are identified, characterise a confounding valley in which holding, spoilage and emission costs cannot be separated from decisions alone, and show the shadow price is recoverable only when the emission signature differs from the cost signature. A numerical study separates finite-sample dispersion from asymptotic bias, documents a bias floor that persists as the sample grows at fixed noise, tests recovery under realistic measurement error and heterogeneous products, validates a bootstrap bias correction, reports confidence interval coverage, and quantifies how often the recovered price is negative. The resulting claims about consistency, identification, and practical reliability are stated in the form that the theory and data jointly support.

Keywords: Inverse optimisation, Deteriorating inventory, Shadow carbon price, Identifiability, Generated regressors, Errors-in-variables, Data driven operations, Sustainable supply chains

1. Introduction

A modern retailer that sells fresh produce, dairy, packaged medicine or any other perishable good records almost everything it does. Every replenishment leaves a timestamp, an order size and a matching stream of sales, and these records accumulate into a dense picture of how the firm actually behaves. What the records never contain is the reason behind the behaviour. We do not see the holding cost the firm assigns to a crowded cold room, we do not see the penalty it attaches to a crate that spoils before it sells, and we certainly do not see the weight, if any, that the firm places on the carbon its operations release. Those quantities live inside a planning system or inside a manager's judgement, and they are exactly the quantities a regulator, an analyst or a rival would most like to know.

The forward direction of inventory theory is mature.

Received: Jul. 29, 2026; Revised: Aug. 28, 2026; Accepted: Oct. 30, 2026; Published: Nov. 13, 2026
Copyright © 2026 Janardan Behera.
DOI: https://doi.org/10.55976/dma.420261668195-237
This is an open-access article distributed under a CC BY license (Creative Commons Attribution 4.0 International License)
https://creativecommons.org/licenses/by/4.0/

Given a cost structure and a demand pattern one computes an optimal replenishment policy, and a century of work beginning with the economic order quantity has refined that computation for deterioration, for finance and for emissions. The inverse direction asks the mirror question. Given the policy a firm has chosen, and given the demand it faced, what cost structure would make that policy optimal. This question is the province of inverse optimisation, a body of technique that has matured in network design, in radiation therapy and in electricity markets but has barely touched the classical inventory models that govern perishable trade. That gap is the opening this paper works in.

Framed this way the exercise is more than an audit. The quantities that behaviour reveals, the deterioration rate and the effective cost-to-holding ratio, are precisely the primitives a planner needs in order to compute future replenishment, so the inverse step is naturally the analytical core of a decision support module rather than a stand-alone diagnostic. We say "module" rather than "system" deliberately: as Section 5.8 makes explicit, what this paper delivers is the estimation engine and a specification of the data it needs, not a deployed, benchmarked, or field-tested software artefact. Read once from the transaction record, the recovered structure can drive the next ordering cycle and flag when a product is drifting away from cost efficient practice. We therefore develop the framework not only as a behavioural lens but as an analytical component that turns the operational record a firm already keeps into forward guidance for inventory planning, the service a decision support system for perishable trade must ultimately provide.

We are drawn to the perishable setting for two reasons. The first is that deterioration makes the inverse problem richer, because the order quantity carries information that pure timing does not. When a firm orders more than it eventually sells, the surplus is spoilage, and the size of that surplus is a direct fingerprint of the deterioration rate, provided, as the revised Section 7 demonstrates, that surplus is not itself contaminated by the very measurement error that afflicts every other quantity in the record. The second reason is topical. Carbon pricing, whether through an explicit levy, a cap-and-trade permit or an internal shadow price, is now a live feature of operational planning, and there is a genuine appetite for knowing whether a firm's replenishment behaviour is consistent with the carbon price it claims to face.

We give that appetite an instrument, but we are careful about what the instrument actually measures. From observed orders and sales, together with the firm's own accounting costs, we recover a model-consistent shadow carbon price: the per unit emission cost under which the firm's observed replenishment is a least cost response, conditional on the forward model being correctly specified, on the accounting costs being accurately known, and on no omitted operational constraint being silently absorbed into the estimate. Section 6 and Section 9 below discuss at length what has to hold for that conditional reading to travel to a claim about a firm's true internal carbon valuation, and what does not.

The technical heart of the paper is a reduction that is almost too tidy. When the cost of a deteriorating item is written per unit time, the fixed and emission linked ordering costs merge into a single effective setup term, and the holding cost, the spoilage cost and the holding linked emission cost merge into a single effective holding term. The replenishment cycle depends on the data only through the ratio of these two effective terms and through the deterioration rate. This collapse is not a modelling convenience, it is the identifiability structure of the problem laid bare. Behaviour can reveal the ratio and the deterioration rate, and nothing more, so any attempt to read an individual cost component out of decisions alone is doomed unless further information is supplied. We turn that obstacle into a method. The deterioration rate comes from the inventory balance, the effective cost-to-holding ratio comes from the first order optimality condition, and the shadow carbon price comes from combining the ratio with the accounting costs the firm already knows.

Because the collapse is so central, we take identifiability seriously as a theoretical object rather than as an afterthought. We prove that the effective ratio and the deterioration rate are the only quantities decisions can reveal, we exhibit a confounding valley of cost structures that are behaviourally indistinguishable, and we derive the precise condition under which the shadow carbon price separates from ordinary cost sensitivity. Section 6 also states explicitly what a single observed record can and cannot identify on its own.

Four questions organise the work. The first (RQ1) asks which features of a firm's cost structure can be recovered from its observed replenishment decisions alone and which are fundamentally hidden. The second (RQ2) asks whether a shadow carbon price, the per unit emission cost a firm behaves as if it faces, can be identified from behaviour together with the firm's accounting costs, and under what condition that recovery is well posed. The third (RQ3) asks how the recovered parameters can be embedded in a decision support workflow so that they guide future inventory planning rather than serving only as a retrospective audit. The fourth (RQ4) asks how the two-stage estimator behaves once its own statistical properties, generated-regressor uncertainty, errors-in-variables, sample heterogeneity, and finite-sample bias at fixed noise, are examined on their own terms rather than assumed away.

The contributions of the paper are the following. We develop a deterministic inverse optimisation framework for deteriorating item inventory under a carbon price, and we show that ordering, holding, spoilage and emission costs collapse into an effective setup and an effective holding pair whose ratio, together with the deterioration rate, exhausts what behaviour can reveal. We construct a two-stage estimator that recovers the deterioration rate

from the inventory balance and the effective cost-to-holding ratio from the optimality condition, and we define a shadow carbon price recovered from the ratio and the accounting costs. We establish an identifiability theory that states what decisions can and cannot reveal, including a confounding valley and an emission signature condition for the carbon price. We characterise the two-stage estimator's finite-sample-noise bias floor analytically and numerically, propose and validate a parametric bootstrap correction, quantify the effect of pooling heterogeneous products, evaluate bootstrap confidence interval coverage, and test recovery once the inventory balance is contaminated by realistic, independent measurement error rather than mechanically preserved by construction. We derive the joint asymptotic distribution of the two-stage estimator via a generalised method of moments (GMM) sandwich formula, propagate accounting-input uncertainty into the shadow carbon price, and implement bias-corrected-and-accelerated and Fieller-type confidence sets for the regime in which ordinary intervals are least trustworthy. We close with a revised sensitivity analysis and a discussion of what the shadow carbon price does and does not license a regulator, an analyst or a firm to conclude.

The remainder of the paper is organised as follows. Section 2 places the work against inverse optimisation, deteriorating inventory, sustainable operations, the econometrics of identification, decision support and data driven decision making, including the statistical treatment of generated regressors and weak identification, and closes with recent developments from 2021 to 2026. Section 3 describes the system, fixes notation, and discusses the identification unit and the conditions under which pooling records across products is defensible. Section 4 develops the forward model and the effective parameter collapse. Section 5 constructs the inverse estimator, discusses its errors-in-variables structure, and sets it inside a decision support workflow. Section 6 proves the identifiability results and adds a treatment of statistical inference. Section 7 reports the numerical study. Section 8 presents the sensitivity analysis. Section 9 discusses managerial and policy implications, and Section 10 concludes with a limitations discussion. Derivations are collected in the appendices.

2. Literature review

2.1 Inverse optimization

Inverse optimisation takes an observed solution as given and searches for objective or constraint parameters that render it optimal. The linear programming foundations were set by Ahuja and Orlin [1], who framed the recovery of a cost vector as a problem of minimal adjustment, and the conic extension followed with Iyengar and Kang [2]. A parallel line grew from convex duality, where Keshavarz et al. [3] imputed an objective from the residual of the stationarity conditions, an idea that sits directly beneath the estimator we use. Once data are noisy the exact rationalisation of every observation becomes impossible, and Aswani et al. [4] established consistency and tractable estimators when observations are corrupted, and Mohajerin Esfahani et al. [5] built a distributionally robust treatment for imperfect information. Online and learning based formulations were developed by Dong et al. [6], and the recovery of constraint parameters rather than objectives was studied by Chan and Kaw [7]. Applications have clustered in energy and pricing, where Saez-Gallego et al. [8] estimated, by inverse optimisation, the price response of a cluster of electricity consumers from historical prices and consumption and used the estimated model to construct market bids. The breadth of the area is surveyed by Chan et al. [9].

2.2 Economic order quantity and deteriorating inventory

The economic order quantity descends from Harris [10] and the dynamic lot sizing tradition from Wagner and Whitin [11], a heritage celebrated a century later in the editorial introduction to a special issue by Cárdenas-Barrón et al. [12]. Deterioration entered the picture with Ghare and Schrader [13], who modelled exponential decay of stock, and with Covert and Philip [14], who allowed a Weibull decay rate. The perishable literature was consolidated by Nahmias [15] and, for the deterministic deteriorating stream specifically, by the reviews of Goyal and Giri [16] and Bakker et al. [17]. Extensions towards finance and credit, such as the time varying deterioration and payment delay of Sarkar [18], show how flexible the forward model has become. Every one of these studies runs in the forward direction. They take the deterioration rate and the cost coefficients as primitives and solve for a policy. We invert that logic and treat the policy as the datum.

2.3 Carbon and sustainable inventory

The coupling of replenishment and emissions is now a substantial field. Hua et al. [19] introduced carbon into the order quantity and traced how a carbon price reshapes the optimal cycle, while Benjaafar et al. [20] embedded emissions in a family of supply chain models to expose the levers a firm actually controls. The carbon constrained order quantity was characterised by Chen et al. [21], and the interaction of replenishment with investment in abatement under several regulatory regimes was studied by Toptal et al. [22]. Sustainable formulations of the order quantity itself appear in Battini et al. [23], and lot sizing under explicit cap and trade or tax regimes in He et al. [24]. Demand that responds to credit under an emission constraint was analysed by Dye and Yang [25]. This work tells us how a known carbon price ought to change

behaviour. The reverse question, of reading an unknown or merely claimed carbon price out of behaviour, has not been posed in this literature, and it is the question this paper addresses.

2.4 Revealed preference and identification

The idea that choices reveal the objective behind them is old and deep in economics. Afriat [26] showed how to construct a utility consistent with observed expenditure, and Varian [27] turned the construction into an operational test of the underlying preferences. Our shadow carbon price is a direct descendant of this tradition, transplanted from the consumer to the replenishing firm. Whether a parameter can be recovered at all is the identification question, whose many meanings across econometrics are catalogued by Lewbel [28]. We borrow that discipline and ask, before estimating anything, what the data are logically capable of revealing. The answer, that decisions pin down a ratio and a decay rate but not the individual costs, is an identification result in exactly the sense that literature intends.

2.5 Estimation from operational data

A growing stream learns primitives from the operational record rather than from surveys or experiments. The dynamic pricing and learning literature surveyed by den Boer [29] learns demand from transactions as they arrive, and the seminal revenue management model of Gallego and van Ryzin [30] shows how tightly price, demand and inventory are linked in a way that any estimator must respect. Our contribution to this stream is to learn cost rather than demand, and to do so from the geometry of an optimality condition rather than from a likelihood.

2.6 Decision support systems and data driven decision making

A decision support system couples data, models and an interface so that a manager acts on analysis rather than on intuition, and the model-driven branch of the field, in which an optimisation or simulation model sits at the core, is the branch our work speaks to [31]. The discipline has long debated its own relevance, research methods and theoretical foundations [32], and the growth of business analytics gave it a unified foundation spanning descriptive, predictive and prescriptive tasks [33]. The prescriptive task, which turns data directly into decisions, has become the centre of gravity, whether by learning a policy that maps observations to actions [34] or by training a predictive model through the very optimisation that will consume it [35]. Both of these assume the objective is known and learn what feeds it. Inverse optimisation occupies the complementary position, learning the objective itself from decisions, and recent methodological advances have carried it into the mixed integer and larger scale settings that operational decision support demands [36], while prescriptive analytics has begun to be applied directly to inventory decisions, for example in hospital restocking [37].

Our contribution sits at this junction: a model driven engine whose distinctive step is to recover the hidden cost objective, including a carbon valuation that a firm's replenishment reveals. As Section 5.8 stresses, however, the paper delivers this engine and a deployment specification, not a benchmarked implementation, so its place in the decision-support literature is that of a candidate model-driven component in the sense of [31], rather than a validated and field-tested decision support system.

2.7 Statistical inverse optimization, generated regressors and weak identification

This subsection situates the paper's statistical machinery within the wider literature on two-stage estimation with generated regressors. The two-stage estimator of Section 5 shares its core statistical difficulty with a wider class of problems: an early stage estimates a nuisance parameter that is then plugged into a later-stage regression, so the later stage's regressors are themselves estimated rather than observed. Hansen [38] establishes the large-sample theory for exactly this class of stacked estimating-equation systems, and we use it directly in the derivation of Section 5.4; Newey [39] and Murphy and Topel [40] give the generated-regressor correction for sequential, two-step estimators that is the template for the bootstrap bias correction we adopt in Section 5.6 and test in Section 7.3, and for the sandwich variance of Section 5.4. Aswani et al. [4], already introduced above, remain the closest existing treatment of inverse optimisation under noisy, imperfectly optimal decision data, a connection we develop further, with a direct simulation comparison, in Section 7.15's joint minimum-distance benchmark under contamination. Newey and McFadden [41] give the general handbook treatment of generated regressors and errors-in-variables in two-stage moment estimators that underlies the formal apparatus we point to in Section 5.6 and use explicitly in Section 5.4, in place of a plain ordinary least squares (OLS) reading, which is insufficient on its own. Fieller [42] gives the classical construction for confidence sets on ratio-valued parameters, precisely the situation our carbon price estimator is in whenever the denominator of Equation (20) is close to zero, and we implement his construction directly, alongside the bias-corrected-and-accelerated (BCa) bootstrap of Efron [43], following the general bootstrap methodology of Efron and Tibshirani [44], in the interval comparison of Section 6.6. On the operations side, Bakker et al. [17] review the deteriorating-inventory literature broadly, classifying the models by the form of the deterioration process and by system features, which is the natural reference point for

extending the cost and emission accounting of Section 5 beyond the single constant rate deterioration channel modelled here.

2.8 Recent developments from 2021 to 2026

Since the foundational contributions reviewed above, inverse optimisation has developed along learning theoretic, algorithmic and applied lines. Besbes et al. [45] study contextual inverse optimisation, characterising a minimax offline policy and a logarithmic regret bound for the online setting, and Zattoni Scroccaro et al. [46] propose an incenter cost vector and an augmented suboptimality loss for inconsistent data, which makes explicit that the choice of loss function is itself a design decision. Gupta and Zhang [47] recover the cost vector of a linear programme from noisy optimal solutions by decomposition and adaptive sampling, and Crönert et al. [48] estimate operational parameters from observed near optimal store locations in a competitive retail location game. The two studies closest to our setting are Huang et al. [49], who apply inverse optimisation to the first order conditions of a buyer's newsvendor problem to infer demand patterns from past transactions, and Erkip [50], who reviews how data can be used to estimate the parameters of richer inventory models. To our knowledge, none of these studies treats a deteriorating item under a carbon price, separates what an observed policy can and cannot identify, or propagates first stage estimation error into a second stage cost ratio, and these are the gaps that the present paper addresses.

On the forward side, carbon regulation has been embedded in deteriorating item models at an accelerating pace. Paul et al. [51] study a green inventory model under a carbon tax with price and green concern dependent demand and deterioration. Jain et al. [52] compare carbon tax, carbon offset and cap and trade regimes in a three-echelon model for deteriorating products and find cap and trade to be the most effective, while Mahato et al. [53] compare cap and price, cap and trade and carbon tax regimes for a perishable retailer with a fixed lifetime and partial backordering. Pervin [54] and Bhadoriya et al. [55] combine deterioration with preservation or green technology investment and emission costs, the former with backorders and the latter with trade credit, shortages and waste disposal. Konstantaras et al. [56] find that the level of a carbon tax materially changes the optimal system cost of a closed loop inventory model. In the cold chain, refrigeration makes holding itself a direct and continuous source of energy use and emissions, so that the holding-linked emission factor becomes a first-order component of the carbon account. All of these studies take the carbon price, or the regulatory regime that sets it, as an input and solve for the policy. Our contribution is the converse step, which asks what carbon price a firm's observed replenishment is consistent with.

The carbon price itself has also attracted recent empirical and conceptual attention. Trinks et al. [57] find that the adoption of internal carbon prices is driven by external carbon constraints and by exposure to formal carbon pricing, which suggests that an internal price and a regulated price are related but distinct. Zakriya et al. [58] report, for a large cross section of firms, that carbon taxes independently improve environmental performance, whereas internal carbon pricing and trading do not when used alone. Zhang et al. [59] show that in a cooperative supply chain the effect of a higher carbon price on retail prices depends on how emission allowances are allocated. These findings motivate a measurable, model consistent shadow price, and they also explain why it must not be equated with a firm's declared internal price without independent evidence, a distinction we maintain throughout.

Recent econometric work bears directly on the inference problems treated in Sections 5 and 6. Hansen and Lee [60] give a variance estimator for iterated generalised method of moments (GMM) estimators that remains consistent under mild misspecification, which supports reporting a sandwich rather than a model-based variance. Dufour et al. [61] use Fieller type confidence sets for ratio defined indices and note that such sets can be unbounded, as in our setting near the identification boundary, and Ratsimalahelo [62] compares delta method, Edgeworth and bias corrected intervals for ratios of parameters with Fieller's interval. Cavaliere et al. [63] show that prepivoting restores the validity of the bootstrap even when the bias cannot be estimated consistently, a result that is relevant to the bootstrap bias correction of Section 5.5 and to the interval comparison of Section 6.6.

Finally, the decision support literature has continued to move towards data driven prescriptive methods. Mandi et al. [64] survey and benchmark decision focused learning, and Sadana et al. [65] unify predict then optimise, decision focused learning and data driven optimisation under the heading of contextual stochastic optimisation. Both lines of work judge an estimator by the quality of the decisions it supports and not by statistical fit alone, which is the standard against which the estimator of Section 5 is positioned in Section 5.8, as a candidate component of a decision support workflow and not as a benchmarked system.

3. Problem description and notation

We observe a collection of replenishment records drawn from a firm or from a set of comparable firms that share a cost structure. "Share a cost structure" is an assumption, not a description, and Section 3.1 below states exactly what it requires and how to check it. A record indexed by i consists of an estimated demand rate D_i, a realised replenishment cycle length T_i and a realised order quantity Q_i. The demand rate is treated as known context, obtained from the sales history, and it varies

across records because products, seasons and price points differ. The cycle length and the order quantity are read directly from the transaction log. The analyst does not observe the cost coefficients, and the central task is to recover as much of them as the records allow.

The forward behaviour of the firm is governed by a deterministic model of a single deteriorating item with continuous demand, which we develop in Section 4. Before that we collect the symbols in Table 1 and record the assumptions that frame the analysis.

Assumption 1 *The firm replenishes a single deteriorating item under a constant deterioration rate, faces a demand rate that is constant within a cycle, holds no safety stock and permits no shortage, so that inventory reaches zero exactly at the end of each cycle.*

Assumption 2 *Each observed cycle length is a least cost response to the demand rate of its record under the common cost structure, possibly perturbed by execution noise or by a bounded tolerance for sub optimality that we make explicit when relevant.*

Assumption 3 *For the recovery of the carbon price the firm's conventional cost coefficients, the ordering cost, the holding cost, the spoilage cost and the two emission intensities, are known to the analyst from accounting records, while the carbon price itself is unknown.*

Assumption 4 *The records pooled into a single estimation exercise are drawn from a single identification unit in the sense of Section 3.1: either repeated cycles of one stock-keeping unit (SKU), or several stock-keeping units that a prior, model-external step has grouped as sharing a common deterioration rate and a common effective cost-to-holding ratio. Pooling records across units that do not satisfy this condition introduces the bias documented in Section 7.8.*

Assumption 1 is the standard deterministic deteriorating item environment. Assumption 2 is the optimality premise that every inverse method requires in some form, and we relax it deliberately in the numerical work to see how recovery survives departures from it. Assumption 3 reflects the realistic asymmetry that a firm knows what a delivery and a cold room cost it in cash but has never made explicit what it charges itself for carbon.

3.1 The identification unit and record pooling

The common-parameter assumption merits closer scrutiny. The estimator of Section 5 recovers a single pair (θ, ρ) from a panel of records, and Theorem 1 is

Table 1. Notation

Symbol	Meaning
D	Demand rate for a record, units per unit time
T	Replenishment cycle length, the decision variable
Q	Order quantity placed at the start of a cycle
$I(t)$	Inventory level at time t within a cycle
θ	Deterioration rate of the item
K	Fixed ordering cost per replenishment
h	Holding cost per unit per unit time
c_d	Spoilage cost per unit that deteriorates
τ	Carbon price per unit of emission
e_o	Emission released per replenishment order
e_h	Emission released per unit held per unit time
A	Effective setup cost, $A = K + \tau e_o$
Ω	Effective holding cost, $\Omega = h + c_d\theta + \tau e_h$
ρ	Effective cost-to-holding ratio, $\rho = \frac{A}{\Omega}$ (units: quantity × time; see section 6.4)
N	Number of replenishment records
$\widehat{(\cdot)}$	Recovered value of a quantity

only a statement about what that panel can identify when the panel truly was generated by one (θ, ρ), i.e. when the records describe one identification unit. Two data-generating settings satisfy this cleanly. The first is repeated replenishment cycles of a single stock-keeping unit observed over time, where the deterioration rate and the cost structure are properties of the physical product and its storage conditions and do not change from cycle to cycle; here demand varies with season or price but the panel is genuinely homogeneous, and the identification result applies as stated, though the empirical range of demand variation will typically be far narrower than the two hundred to three thousand unit range used for illustration in Section 7. The second is a designed cross-section of several stock-keeping units that a product taxonomy, a shelf-life specification, or a prior clustering step has already grouped as physically and operationally comparable, for example, several SKUs of the same packaged salad line stored in the same cold chain.

Pooling records across SKUs whose deterioration rates differ materially, different produce categories, different refrigeration regimes, different shelf lives, is a misspecification, not a robustness question, and Section 7.8 shows what it costs: a pooled estimator can converge to a precise but wrong common θ and ρ, with sampling variance shrinking exactly as if the pooling were correct while the point estimate stays biased. We recommend, and Section 7.8 evaluates, three practical responses that require no change to the core estimator: (i) fit the two-stage estimator separately within each product cluster and report cluster-specific (θ, ρ, τ) rather than a single pooled figure; (ii) where clusters are not known a priori, use a preliminary clustering step on observed spoilage-to-order ratios, which are informative about θ even before the formal estimator is run, to propose candidate groups; and (iii) treat a materially rising within-cluster optimality-fit residual, tracked over Section 5.6's benchmarking workflow, as a signal that a cluster is not homogeneous and should be split. A fully specified random-coefficients or hierarchical extension that estimates a distribution of (θ, ρ) across products jointly, rather than a single pooled value or a set of disjoint cluster values, is a natural next step and is flagged as future work in Section 10.3; it is a materially larger estimation problem than the one this paper solves and we do not attempt it here.

This identification-unit requirement raises a design question the numerical study had not directly addressed: every experiment above draws D_i from a wide range (200 to 3,000 units) that better represents variation across different stock-keeping units than the narrower variation a single SKU exhibits from one replenishment cycle to the next. We repeated the clean-design and contaminated-design experiments with D_i drawn instead from narrow within-SKU ranges centred on 1,000 units, from ± 30% down to ± 5%, and separately varied the panel size at a fixed ± 10% range to reflect realistic annual cycle counts for a single product. Table 2 reports the resulting root mean squared errors (RMSE).

On the mechanically exact design, narrowing the demand support has essentially no effect on $\hat{\theta}$, which continues to recover exactly regardless of how much curvature D_i contributes, and if anything a mildly favourable effect on RMSE($\hat{\rho}$), which falls slightly as the range narrows from ± 30% to ± 10% before flattening. This is a more reassuring finding than the identification-unit argument alone would have predicted, and the reason is that the second-stage regression is through the origin: what identifies ρ is the average ratio x_{1i}/x_{2i} weighted by x_{2i}^2, not the cross-record spread in x_{2i} itself, so a within-SKU panel is not disadvantaged relative to a cross-SKU one merely for having less demand heterogeneity. What does matter, unsurprisingly, is the number of cycles: at a fixed ± 10% within-SKU range, RMSE($\hat{\rho}$) fell from 0.220 at 20 cycles to 0.139 at 50, 0.114 at 100 and 0.108 at 200, the familiar root-N pattern. Under contamination the picture is more mixed: RMSE($\hat{\theta}$) is marginally lower for the narrow within-SKU design than for the cross-SKU baseline (0.362 against 0.433), but RMSE($\hat{\rho}$) is somewhat higher (0.531 against 0.445), an approximately twenty per cent increase consistent with weaker identifying variation becoming a second-order factor once the balance equation is no longer exact. We conclude that the identification-unit requirement of Assumption 4 remains a correct and necessary modelling discipline, pooling across genuinely

Table 2. Effect of narrowing the demand support from cross-SKU to within-SKU variation, $N = 200$ unless noted. Contaminated rows use 2% order-quantity noise plus 3% ending inventory. "—" indicates not evaluated at that width for the contaminated design

Demand-support design	RMSE($\hat{\theta}$), clean	RMSE($\hat{\rho}$), clean	RMSE($\hat{\theta}$), contam.	RMSE($\hat{\rho}$), contam.
Cross-SKU, 200–3,000 (baseline)	0.000	0.121	0.433	0.445
Within-SKU ± 30% (700–1,300)	0.000	0.109	—	—
Within-SKU ± 20% (800–1,200)	0.000	0.104	—	—
Within-SKU ± 10% (900–1,100)	0.000	0.102	0.362	0.531
Within-SKU ± 5% (950–1,050)	0.000	0.103	—	—

heterogeneous units still introduces the bias documented in Section 7.8, but that the concern that a realistically narrow within-SKU demand range would by itself starve the estimator of identifying variation is not borne out by direct simulation: sample size and balance-equation reliability are both larger determinants of precision than demand-range width.

4. Forward model

Within a cycle the inventory depletes through demand and through deterioration, in the classical exponential-decay tradition originating with Ghare and Schrader [13], which gives the linear dynamics

$$\frac{dI(t)}{dt} = -D - \theta I(t), \quad 0 \le t \le T, \quad T(t) = 0. \tag{1}$$

Integrating (1) backward from the terminal condition yields the inventory path and the order quantity,

$$I(t) = \frac{D}{\theta}\left[e^{\theta(T-t)} - 1\right], \quad Q = I(0) = \frac{D}{\theta}\left[e^{\theta T} - 1\right]. \tag{2}$$

The time integral of stock over a cycle and the quantity that spoils follow as

$$\int_0^T I(t)\,dt = \frac{D}{\theta^2}\left[e^{\theta T} - 1 - \theta T\right], \quad Q - DT = \theta \int_0^T I(t)\,dt. \tag{3}$$

The detailed derivations of the inventory path, order quantity, holding integral and spoilage identity are provided in Appendix A. The final identity in (3) is worth pausing on. The units that spoil are exactly the deterioration rate multiplied by the holding integral, which means spoilage and holding are not independent cost drivers but two readings of the same stock exposure. This is the first hint of the collapse that organises the whole paper.

The firm bears a fixed ordering cost, a holding cost on the stock integral, a spoilage cost on the units that decay and a carbon charge on the emissions from ordering and from holding. Writing the cost per unit time and dividing the cycle cost by T gives

$$C(T) = \frac{K + \tau e_o}{T} + (h + \tau e_h) D \widehat{H}_b(T;\theta) + c_d \theta D \widehat{H}_b(T;\theta), \tag{4}$$

$$\widehat{H}_b(T;\theta) = \frac{e^{\theta T} - 1 - \theta T}{\theta^2 T}, \tag{5}$$

where $\widehat{H}_b$ is the holding integral per unit time with the demand rate factored out. Collecting terms, the cost rate reduces to a two-parameter form,

$$C(T) = \frac{A}{T} + \Omega D \widehat{H}_b(T;\theta), \quad A = K + \tau e_o,$$

$$\Omega = h + c_d \theta + \tau e_h. \tag{6}$$

The fixed ordering cost and the ordering emission charge merge into an effective setup A. The holding cost, the spoilage cost and the holding emission charge merge into an effective holding Ω. The cycle length that the firm chooses is determined by (6), and therefore by A, Ω, and θ alone. No other feature of the cost structure influences the decision.

The optimal cycle solves the first order condition obtained by differentiating (6),

$$\frac{\partial C}{\partial T} = -\frac{A}{T^2} + \Omega D \widehat{H}_b^{\prime(T;\theta)} = 0. \tag{7}$$

The second derivative of C is positive throughout the feasible region because the holding integral is convex and increasing in T, so the stationary point is the unique global minimiser. Rearranging (7) isolates the object that behaviour will later reveal,

$$A = \Omega D T^{*2} \widehat{H}_b^{\prime(T^*;\theta)}, \quad \text{equivalently} \quad D \widehat{H}_b^{\prime(T^*;\theta)} = \frac{\rho}{T^{*2}},$$

$$\rho = \frac{A}{\Omega}. \tag{8}$$

As the deterioration rate vanishes the model returns to the classical order quantity: the optimal cycle collapses to the square root formula with A in the role of the setup cost and Ω in the role of the holding cost. The deteriorating model is thus a smooth generalisation of the order quantity in which the two classical costs are replaced by their effective counterparts. The detailed mathematical derivations of the effective setup and holding costs, together with the uniqueness of the optimal cycle, are presented in Appendix B.

4.1 Scope of the cost and balance structure

Two simplifications in the model above deserve explicit acknowledgement rather than remaining implicit. First, the emission structure carries only two channels, a fixed emission per order, e_o, and a rate of emission per unit held per unit time, e_h, and excludes deterioration-disposal emissions, nonlinear refrigeration energy use that would make e_h depend on ambient conditions or fill level, transport-distance effects, shipment-size economies, and reprocessing emissions. We keep the two-channel structure because it is the minimal specification for which Theorem 2's identification condition has a clean statement, and because, as Section 9 notes, cold-chain holding emissions are frequently the dominant and most

nonlinearly-behaved channel in practice, which the e_h term is built to capture at least to first order. A richer specification, for example temperature- and age-dependent deterioration and disposal emissions, building on the classification of deterioration forms in Bakker et al. [17] for the age dependent part and on cold chain energy and emission studies such as Hu et al. [66] for the temperature dependent part, is a natural extension we do not carry out here; Section 10.3 flags it as future work rather than folding in an untested extension.

Second, Assumption 1's no-shortage, no-safety-stock, zero-ending-inventory structure, under which the entire gap between order quantity and sales is attributed to deterioration, is restrictive relative to most real perishable-inventory systems, which also contend with returns, shrinkage, physical damage, donation, theft, and stock carried across cycles. We do not relax Assumption 1 here: doing so correctly would change the forward model of this section and every downstream derivation, which is a substantial undertaking in its own right. The numerical study does not imply that Assumption 1 is innocuous. Section 7.5 introduces exactly this kind of contamination, independent order-quantity measurement error and nonzero ending inventory, directly into the data-generating process and shows how much of Section 5.1's deterioration-rate accuracy depends on Assumption 1 holding closely; a fuller treatment that models returns, shrinkage, donation and theft as separate, estimable channels rather than as generic noise is future work (Section 10.3).

5. Inverse optimization framework

The analyst holds the records $\{(D_i, T_i, Q_i)\}$ for $i = 1, ..., N$ and seeks the cost structure behind them. The forward model tells us at once that this cannot be done component by component, because decisions depend on the costs only through A, Ω, and θ. We therefore recover what is recoverable in two stages and then convert the recovered quantities into the carbon price under Assumption 3.

5.1 Recovering the deterioration rate

The deterioration rate is written into the relationship between what is ordered and what is sold. The inventory balance (2) states that the order quantity equals $\left(\frac{D}{\theta}\right)(e^{\theta T} - 1)$, a relationship that holds by the physics of decay and owes nothing to optimality. Given the records the rate is the minimizer of the balance residual.

$$\hat{\theta} = \arg\min_{\theta > 0} \sum_{i=1}^{N} \left[Q_i - \frac{D_i}{\theta}(e^{\theta T_i} - 1) \right]^2. \tag{9}$$

This channel is strong precisely when deterioration matters, because the surplus of the order over sales grows with the rate. It is also only as strong as the balance equation is exact. Equation (9) treats every unit of the observed gap $Q_i - D_i T_i$ as spoilage, which is true by construction whenever Q_i is generated from the disturbed cycle length through the same balance equation, as in the numerical design of Section 7. It is not true when the order quantity itself is measured with independent error, when ending inventory is not exactly zero, when returns or shrinkage are present, or when package sizes and minimum order quantities round the order away from its balance-implied value. Section 5.6 below and Section 7.3 quantify how much this matters.

5.2 Recovering the effective cost-to-holding ratio

With the deterioration rate in hand, the effective cost-to-holding ratio follows from the optimality condition (8). Define for each record the two constructed quantities

$$x_{1i} = D_i \hat{H}_b'^{(T_i;\hat{\theta})}, \quad x_{2i} = \frac{1}{T_i^2}. \tag{10}$$

At a least cost decision the condition (8) requires $x_{1i} = \rho x_{2i}$, so the ratio is the slope of a regression through the origin of x_1 on x_2. Its ordinary least squares estimate is closed form,

$$\hat{\rho} = \frac{\sum_i x_{1i} x_{2i}}{\sum_i x_{2i}^2}. \tag{11}$$

5.3 What kind of regression this actually is

Treating (11) as an ordinary least squares fit and invoking the law of large numbers and the central limit theorem without further comment is not adequate, for three reasons that we now state explicitly rather than leave implicit.

First, both constructed regressors in (10) are nonlinear, correlated functions of the same noisy cycle length T_i: x_{1i} depends on T_i through $\hat{H}_b'(T_i;\hat{\theta})$, and $x_{2i} = 1/T_i^2$ depends on it directly. This is not a design in which the regressor is fixed or measured without error and the noise enters only through the dependent variable; it is closer to an errors-in-variables regression in which the same measurement error contaminates both sides of the equation through a common latent variable. Second, x_{1i} is built from $\hat{\theta}$, a first-stage estimate rather than the true θ, so the second-stage regression carries generated-regressor uncertainty: the sampling variability of $\hat{\theta}$ propagates into x_{1i} and, through it, into $\hat{\rho}$, in a way that a textbook OLS variance formula for (11) does not capture. Third, and most consequentially, the regressors are nonlinear transformations of a noisy T_i, so even at the true θ, taking expectations term by term

does not commute with the nonlinear map; the second-order curvature term derived in Appendix D is exactly the algebraic signature of this non-commutativity, and it is what produces a bias that does not vanish as N grows at fixed noise, see the revised Proposition 4 and Section 7.3.

A fully rigorous treatment casts the two stages as a stacked set of estimating equations and derives the joint asymptotic distribution of (θ, ρ, τ) by the standard GMM sandwich formula, with the first-stage Jacobian propagated into the second-stage variance through the delta method, following the general apparatus of Hansen [38], Newey [39] and Murphy and Topel [40]. We carry out that derivation in full in Section 5.4 immediately below, rather than deferring it: it stacks the two normal equations into a single estimating-equation system, derives the block-triangular Jacobian in closed form, and validates the resulting sandwich variance numerically against Monte Carlo simulation both on the mechanically exact design and under the realistic contamination of Section 7.5. What that derivation provides, and what a working analyst needs most immediately, is (i) an explicit, closed-form account of why and by how much (11) understates its own uncertainty once the balance equation is contaminated, (ii) the numerical characterisation of how large the resulting bias and under-coverage actually are (Sections 7.3 and 7.10), and (iii) a simple, generic bias correction, formally justified rather than only offered as a practical stopgap, that a practitioner can apply today.

5.4 The stacked estimating equations and the GMM sandwich variance

We now carry out the derivation Section 5.3 previously deferred. Write the stage-1 normal equation implicit in (9) and the stage-2 normal equation implicit in (11) as a single stacked system of estimating equations in (θ, ρ). For record i define

$$\psi_{1i}(\theta) = [Q_i - g(D_i, T_i; \theta)] \cdot g'^{(D_i, T_i; \theta)}, \tag{12}$$

which is the derivative, set to zero and summed, of the stage-1 sum of squares in (9), where $g(D, T; \theta) \equiv \left(\frac{D}{\theta}\right)(e^{\theta T} - 1)$ and g' denotes $\frac{\partial g}{\partial \theta}$, and

$$\psi_{2i}(\theta, \rho) = x_{2i} \cdot [x_{1i}(\theta) - \rho x_{2i}], \tag{13}$$

which is likewise the derivative of the stage-2 sum of squares behind (11), with $x_{1i}(\theta)$ and x_{2i} as defined in (10). By construction $\hat{\theta}$ and $\hat{\rho}$ jointly solve $\sum_i \psi_{1i}(\theta) = 0$ and $\sum_i \psi_{2i}(\theta, \rho) = 0$, so ($\hat{\theta}$, $\hat{\rho}$) is exactly the two-parameter, two-equation Z-estimator whose asymptotic distribution a fully rigorous treatment of the two-stage estimator requires. Standard M-estimation theory (Hansen [38], Newey [39], Murphy and Topel [40], Newey and McFadden [41]) gives

$$\sqrt{N}(\hat{\theta} - \theta_0, \hat{\rho} - \rho_0) \to N\ (0, J^{-1} \Sigma J^{-\top}), \tag{14}$$

where $J = \mathrm{E}\left[\frac{\partial \psi}{\partial(\theta, \rho)}\right]$ is the population Jacobian of the stacked moments and $\Sigma = \mathrm{Var}(\psi_i)$ is the covariance of the per-record moment contributions, with $\psi_i = (\psi_{1i}, \psi_{2i})^{\top}$. Because ψ_{1i} does not involve ρ, J is block lower triangular,

$$J_{11} = \mathrm{E}[-g'^2 + (Q_i - g)g''],\ J_{21} = \mathrm{E}\left[x_{2i} \cdot D_i \cdot \frac{\partial \hat{H}_b'}{\partial \theta}\right],$$

$$J_{22} = -\mathrm{E}[x_{2i}^2] \tag{15}$$

evaluated at (θ_0, ρ_0). On the mechanically exact balance design used throughout Section 7, $Q_i = g(D_i, T_i; \theta_0)$ holds identically for every record, so the second term in J_{11} vanishes not merely in expectation but pointwise, and $\psi_{1i}(\theta_0) \equiv 0$ for every i: the first stage has literally zero sampling variance, which is exactly why $\hat{\theta}$ recovers θ_0 to machine precision throughout Sections 7.2, 7.3, and 7.8 regardless of noise, sample size, or perishability. J_{21}, the channel through which first-stage uncertainty would ordinarily propagate into the second stage as a generated-regressor correction, is then multiplied by a first-stage variance of zero and drops out. Once the balance equation is contaminated as in Section 7.5, $\mathrm{Var}(\psi_{1i})$ is no longer zero, J_{21} stays nonzero, and the generated-regressor correction becomes the dominant term in the sandwich, exactly tracking the explosive growth in the empirical RMSE documented there.

Inverting the lower-triangular J and reading off the two diagonal blocks and the cross term of $J^{-1}\Sigma J^{-\top}$ gives closed-form large-sample variances

$$\mathrm{Var}(\hat{\theta}) \approx \frac{\Sigma_{11}}{N J_{11}^2}, \tag{16}$$

$$\mathrm{Var}(\hat{\rho}) \approx \frac{1}{N}\left[\frac{J_{21}^2 \Sigma_{11}}{J_{11}^2 J_{22}^2} - 2\frac{J_{21}\Sigma_{12}}{J_{11} J_{22}^2} + \frac{\Sigma_{22}}{J_{22}^2}\right] \tag{17}$$

with the corresponding covariance

$$\mathrm{Cov}(\hat{\theta}, \hat{\rho}) \approx \frac{1}{N}\left[-\frac{J_{21}\Sigma_{11}}{J_{11}^2 J_{22}} + \frac{\Sigma_{12}}{J_{11} J_{22}}\right]. \tag{18}$$

A further delta-method step through $\tau = f(\rho, \theta; K, h, c_d, e_o, e_h)$ in Equation (20) gives

$$\mathrm{Var}(\hat{\tau}) \approx \left(\frac{\partial \tau}{\partial \rho}\right)^2 \mathrm{Var}(\hat{\rho}) + \left(\frac{\partial \tau}{\partial \theta}\right)^2 \mathrm{Var}(\hat{\theta})$$

$$+2\frac{\partial\tau}{\partial\rho}\frac{\partial\tau}{\partial\theta}\mathrm{Cov}(\hat{\theta},\hat{\rho}) \quad (19)$$

with the partial derivatives given in closed form in Section 5.7 below, where they are combined with the analogous partial derivatives with respect to the accounting inputs K, h, c_d, e_o, and e_h.

We evaluated (15)–(17) numerically (population moments and their Jacobian estimated from a large N = 20,000 reference panel, so that Monte Carlo noise in J and Σ themselves is negligible) and compared the resulting asymptotic standard errors against the empirical standard deviation of $\hat{\theta}$, $\hat{\rho}$, and $\hat{\tau}$ over 400 independently generated panels of the paper's baseline size N = 200. Table 3 reports the comparison for the mechanically exact design of Section 7.2 and for the realistic contamination design of Section 7.5.

The sandwich formula tracks the Monte Carlo standard deviation closely under both designs, including correctly predicting zero first-stage variance on the mechanically exact design and a nonzero, materially larger variance once the balance equation is contaminated. Two qualifications are appropriate. First, (16)–(17) are first-order (variance-only) large-sample approximations; they do not, and are not intended to, capture the second-order bias term of Appendix D and the revised Proposition 4, which is a separate, deterministic component of RMSE and is already reported alongside sampling variance in Section 7.3. The sandwich variance and the bias floor are complements, not substitutes, and both should be read together, exactly as Section 7.4 already does for $\hat{\tau}$. Second, at N = 200 the discrepancy between the theoretical and Monte Carlo standard errors under contamination (2.84 against 2.31 for $\hat{\tau}$) is larger than under the clean design, reflecting the usual higher-order finite-sample slippage of a first-order asymptotic approximation once the estimating equations are less smooth in the contaminating noise; the approximation improves as N grows, consistent with its asymptotic justification. Subject to these two qualifications, (14)–(17) are a defensible large-sample variance expression for the two-stage estimator, derived under the clearly stated asymptotic regime of a correctly specified stacked estimating-equation system with i.i.d. records and the Jacobian evaluated at the truth.

5.5 A parametric bootstrap bias correction

The correction is a one-step parametric bootstrap in the spirit of a jackknife bias correction. Having obtained $(\hat{\theta}, \hat{\rho})$ from the original panel, reconstruct the cycle length that $(\hat{\theta}, \hat{\rho})$ implies for every observed demand rate, re-attach the same noise process that generated the data, the same execution-noise level, resampled a large number of times, refit $(\hat{\theta}, \hat{\rho})$ on each synthetic panel, and take the average discrepancy between the refitted $\hat{\rho}$ and the value $\hat{\rho}$ used to generate the synthetic panels as an estimate of the noise-induced bias. Subtracting that estimated bias from the original $\hat{\rho}$ gives a bias-corrected estimator $\hat{\rho}_{\mathrm{bc}}$. Section 7.3 reports that, at the paper's baseline design (N = 200, three per cent execution noise), this correction reduces the bias in $\hat{\rho}$ from –0.083 to –0.005 and the RMSE from 0.117 to 0.083, essentially removing the systematic component and leaving only sampling noise. The correction requires no change to the estimator itself, only repeated re-simulation around the fitted parameters, and it is the practical answer we offer in place of a full closed-form asymptotic bias formula.

5.6 Recovering the shadow carbon price

The ratio and the rate are all that behaviour reveals, but they are enough to expose the carbon price once the conventional costs are known. Substituting the definitions of A and Ω into $\rho = \frac{A}{\Omega}$ and solving for the carbon price gives

$$\hat{\tau} = \frac{K - \hat{\rho}(h + c_d\hat{\theta})}{\hat{\rho}e_h - e_o}. \quad (20)$$

Table 3. Sandwich-formula asymptotic standard errors (SE) against Monte Carlo standard deviations (SD), mechanically exact and contaminated designs, baseline cost structure, N = 200

Design	Quantity	Sandwich formula SE	Monte Carlo SE (400 reps)
Mechanically exact balance	$\hat{\theta}$	0.0000	0.0000
Mechanically exact balance	$\hat{\rho}$	0.0834	0.0901
Mechanically exact balance	$\hat{\tau}$	0.588	0.648
2% order-qty noise + 3% ending inv.	$\hat{\theta}$	0.0273	0.0240
2% order-qty noise + 3% ending inv.	$\hat{\rho}$	0.0920	0.0968
2% order-qty noise + 3% ending inv.	$\hat{\tau}$	2.838	2.311

Table 4. Delta-method variance decomposition of $\hat{\tau}$ with respect to the five accounting inputs, baseline cost structure

Input	Assumed rel. SD	$\frac{\partial \tau}{\partial(\text{input})}$	Variance contribution	Share
K	5%	0.415	17.19	34.7%
h	5%	–6.097	2.32	4.7%
c_d	8%	–1.829	4.82	9.7%
e_o	10%	10.365	4.30	8.7%
e_h	10%	–152.43	20.91	42.2%

We call τ the shadow carbon price, for the reasons given in Section 1 and returned to in Section 6.4 and Section 9. It is the per unit emission cost under which the firm's observed replenishment is a least cost response given everything else the firm is known to spend, and it is a model-conditional quantity: Assumption 3's accounting inputs, Assumption 1's structural form, and Assumption 2's optimality premise all have to hold for τ to mean what Equation (20) says it means. A firm that ignores carbon will produce a shadow price near zero, and a firm that prices carbon aggressively will produce a large one. The denominator in (20) carries a warning that we develop formally in the next section: when it is small the recovered price is sensitive to any error in the ratio, and when it vanishes the price is not identified at all. Because τ is a ratio-type statistic near a potentially small denominator, it can also take an economically inadmissible negative value under noise even when the true price is non-negative; Section 7.10 quantifies how often this happens and evaluates truncating τ at zero as a simple remedy, in place of leaving the raw estimate unconstrained. The mathematical details of the two-stage estimation procedure and the derivation of the shadow carbon price are provided in Appendix C.

5.7 Propagating uncertainty in the accounting inputs

Equation (20) treats K, h, c_d, e_o, and e_h as known constants, but Assumption 3 only requires that the analyst have accounting-quality estimates of them, not that those estimates be exact. Measurement error in these five inputs, rather than sampling noise in $\hat{\rho}$ and $\hat{\theta}$, may in fact be the binding constraint on how precisely $\hat{\tau}$ can be known in practice, and this subsection quantifies that concern rather than leaving it as a caveat. Differentiating (20), and writing $D \equiv \rho e_h - e_o$ for the identification denominator, gives the closed-form sensitivities

$$\frac{\partial \tau}{\partial K} = \frac{1}{D}, \quad \frac{\partial \tau}{\partial h} = -\frac{\rho}{D}, \quad \frac{\partial \tau}{\partial c_d} = -\frac{\rho\theta}{D},$$

$$\frac{\partial \tau}{\partial e_o} = \frac{\tau}{D}, \quad \frac{\partial \tau}{\partial e_h} = -\frac{\tau\rho}{D}. \tag{21}$$

A first-order (delta-method) variance decomposition is

$$\mathrm{Var}(\hat{\tau}) \approx \sum_k \left(\frac{\partial \tau}{\partial k}\right)^2 \mathrm{Var}(k), \quad k \in \{K, h, c_d, e_o, e_h\} \tag{22}$$

holding $\hat{\rho}$ and $\hat{\theta}$ at their point estimates; the full expression combining this input-uncertainty term with the sampling variance of $\hat{\rho}$ and $\hat{\theta}$ from Section 5.4 is

$$\mathrm{Var}(\hat{\tau})_{\text{total}} \approx \mathrm{Var}(\hat{\tau})_{\text{sampling}} + \mathrm{Var}(\hat{\tau})_{\text{input}} \tag{23}$$

the two being independent because the accounting inputs are measured separately from the transaction log records. At the baseline cost structure (K = 200, h = 5, c_d = 15, θ = 0.30, e_o = 2.0, e_h = 0.30, ρ = 14.71), and assuming relative measurement error of 5% on K and h, 8% on c_d (spoilage cost is typically the hardest of the five to cost precisely), and 10% on the two emission intensities e_o and e_h (consistent with the uncertainty ranges reported for activity-based carbon accounting), Table 4 reports the resulting variance decomposition.

Two findings follow. First, at these plausible accounting error levels, input uncertainty contributes a standard deviation of roughly 7.0 units to $\hat{\tau}$, an order of magnitude larger than the sampling standard deviation of 0.4–0.5 units reported for the same design in Section 7.3. In this cost structure, the dominant source of uncertainty in the recovered shadow carbon price is not how many replenishment records the analyst has, but how precisely the firm's own accounting department can state its holding, ordering and spoilage costs and its two emission intensities. Second, the sensitivities are highly uneven: e_h alone accounts for 42% of the input-uncertainty variance, followed by K at 35%, while h, c_d, and e_o together account for the remaining 23%. This traces directly to the amplification factor $\frac{\partial \tau}{\partial e_h} = -\frac{\tau\rho}{D}$, which grows without bound as the identification denominator $D = \rho e_h - e_o$

shrinks, exactly the conditioning problem already central to Proposition 3 and Section 7.9. A practical implication for deployment is that measurement effort is better spent tightening the estimate of e_h, the emission intensity of holding, than on any of the four conventional cost inputs, whenever the firm's cost structure places it near the identification boundary. We combined the two sources of uncertainty in a joint Monte Carlo experiment (N = 200 sampling and independent lognormal perturbation of the five inputs at the relative errors of Table 4, 2,000 replications): the resulting 90% interval for $\hat{\tau}$ was [15.6,40.3] (standard deviation 7.78), against [24.5,26.6] (standard deviation 0.61) when accounting inputs are held fixed at their true values and only sampling variation operates, confirming that the two sources are of very different order and that reporting a confidence interval for $\hat{\tau}$ without accounting for input uncertainty materially overstates the precision with which the shadow carbon price is known.

5.8 Decision support workflow

The recovered quantities are not an end in themselves. They are the objective a forward inventory model needs in order to prescribe decisions, so the inverse step is naturally the analytical core of a decision support workflow for perishable replenishment. We describe this as a workflow rather than a decision support "system", because the paper specifies the computation and its data requirements but does not implement, benchmark, or field-test a system around it. The workflow reads order sizes, sales rates, realised cycle lengths and accounting costs from the transaction and enterprise resource planning records a firm already maintains, applies the two stage recovery of this section, and the recovered structure then feeds three uses: forward replenishment planning, since the effective cost-to-holding ratio and the deterioration rate are exactly the primitives Equation (8)'s forward optimality condition needs, without requiring the effective setup and the effective holding separately; monitoring the shadow carbon price period by period as a running indicator rather than a one-off audit; and benchmarking and drift detection, since the residual of the optimality fit measures how far behaviour departs from any coherent cost structure. Section 5.9 below adds a diagnostic threshold for that residual.

Table 5 below is a compact reference for practitioners deciding what to report and when to withhold a number rather than report it with false confidence.

5.9 A practical threshold for the optimality-fit residual

The residual sum of squares in the second-stage fit, $\sum_i (x_{1i} - \rho x_{2i})^2$, is proposed as a diagnostic of Assumption 2, but scaling the residual by a chi-squared reference with $N - 1$ degrees of freedom, as an earlier draft of this subsection recommended as the primary rule, is only as reliable as the variance estimate it is scaled by, and Table 6 below shows that rule has poor size control in practice. We therefore reverse the priority given to the two rules. The primary recommendation is now the bootstrap alternative: resample records under the fitted (θ, ρ) and the

Table 5. Roadmap for the three recovered quantities: data requirements, identification condition, principal bias source, recommended uncertainty treatment, and circumstances warranting withholding the estimate

Output	Data req.	Identification condition	Principal bias source	Recommended uncertainty method	Do not report when
θ	D_i, T_i, Q_i per record	Balance Eq. (9) exact or near-exact; Assumptions 1–2	Order-qty/ending-inv. contamination (§7.5); negligible if balance exact	Sandwich variance (§5.4); bootstrap-calibrated residual check (Table 6)	Balance-equation reliability cannot be verified: supply θ externally instead (§7.5)
ρ	$\hat{\theta}$ (or external θ), T_i per record	Identified up to the ratio only (Prop. 1–2); single identification unit (Assumption 4)	Generated-regressor propagation from $\hat{\theta}$ (§5.4); pooling heterogeneous units (§7.7)	Sandwich variance + bias correction (§5.5); BCa/ Fieller-type interval (§6.6)	Records pooled across units violating Assumption 4; conditioning denominator not distinguishable from zero
τ	$\hat{\rho}$, $\hat{\theta}$, externally supplied K, h, c_d, e_o, e_h	Identification denominator $\rho e_h - e_o$ bounded away from zero (Prop. 3)	Accounting/emission input measurement error (§5.7), typically dominant over sampling error	Combined sampling + input-uncertainty variance (§5.4 + §5.7); one-sided interval near $\tau = 0$ (§7.9); Fieller-type set near the boundary	Input uncertainty in e_h or K unquantified; denominator near zero without a Fieller-type caveat; raw negative estimate reported without one-sided reframing

Table 6. Empirical false alarm rate (null row) and power (remaining rows) of the chi-squared and bootstrap-calibrated residual diagnostics, $N = 200$, 150 independent panels per regime, nominal levels 95% and 99%

Regime	Chi-sq. rule @95%	Chi-sq. rule @99%	Bootstrap rule @95%
Null: Assumption 2 holds exactly (nominal 5%/1%)	7.3%	3.3%	3.3%
Parameter heterogeneity across pooled products	100.0%	100.0%	100.0%
Bounded rationality, 3% tolerance	74.0%	62.0%	74.7%
Balance-equation contamination (§7.5)	31.3%	21.3%	6.7%

estimated noise level, build the empirical distribution of the residual sum of squares under the null that Assumption 2 holds, and flag a panel when its observed residual falls in the upper tail of that empirical distribution, using the same bootstrap machinery already used for the bias correction in Section 5.5 and the coverage evaluation in Section 7.11, so no new computational infrastructure is required. The chi-squared rule, scaled by a residual variance estimated from a trusted low-noise reference period for the same product where one exists, is retained only as a quick, back-of-envelope check, with the explicit caveat established below that it distorts the nominal false alarm rate, sometimes substantially. Table 6 reports the empirical false alarm rate of both rules under the null (Assumption 2 satisfied exactly) and their power against three specific violations: parameter heterogeneity across pooled products, bounded rationality with a three per cent tolerance around the least cost cycle, and the balance-equation contamination of Section 7.5.

The chi-squared rule over-rejects under the null at both nominal levels; the distortion is worst at the 99% level, where it flags more than three times as many compliant panels as it should (3.3% observed against a 1% nominal target). The bootstrap rule controls size correctly at the 95% level by construction. Against heterogeneity and bounded rationality both rules have comparable, high power. Against balance-equation contamination the two rules diverge sharply: the bootstrap rule has markedly lower power (6.7%) than the chi-squared rule (31.3%). The reason is instructive rather than a flaw to be hidden. Balance contamination biases the fitted $\hat{\theta}$ itself, and the bootstrap threshold is calibrated by resampling around that same biased $\hat{\theta}$, so part of the contamination is absorbed into the null distribution the panel is judged against. The chi-squared rule, calibrated externally from a trusted reference period, is not subject to this self-referential blind spot and is, on this one failure mode only, the more sensitive of the two. We therefore recommend the bootstrap rule as the primary, correctly sized general-purpose diagnostic, and recommend pairing it with the direct balance-equation check of Section 7.5 (comparing observed order quantities against $D_i \cdot T_i$) specifically to catch the contamination pattern the bootstrap residual is least able to see.

6. Identifiability

The forward model already signals that behaviour cannot reveal every cost. This section makes the claim precise. We state what a body of replenishment records can identify, we exhibit the exact indeterminacy that remains, and we settle when the shadow carbon price is meaningful. Unless stated otherwise, records in this section are generated by a common cost structure under Assumption 1 and Assumption 2 (the exact, noiseless case); Section 6.5 and Section 7 examine what changes once noise, generated-regressor uncertainty and heterogeneity are introduced.

6.1 Effective sufficiency and scale indeterminacy

Proposition 1 (Effective sufficiency). *The optimal cycle length depends on the cost structure only through the triple* (A, Ω, θ). *Two structures that share this triple generate identical replenishment for every demand rate.*

Proof. The cost rate (6) depends on the coefficients only through A, Ω, and θ, and the minimiser of a function is determined by the function. Any two structures with a common triple therefore share the objective $C(\cdot)$ at every demand rate and share its minimiser. □

Proposition 2 (Scale indeterminacy). *The optimal cycle length is invariant to the rescaling* $(A, \Omega) \to (\lambda A, \lambda\Omega)$ *for any* $\lambda > 0$. *Consequently the level of costs is not identified from decisions, and only the ratio* $\rho = \dfrac{A}{\Omega}$ *and the rate* θ *can be.*

Proof. Under the rescaling the objective becomes $\lambda C(\cdot)$, whose minimiser is unchanged. Decisions are thus constant along every ray of fixed ρ, so no function of the data can separate A from Ω beyond their ratio. □

6.2 What a single record identifies

A natural conjecture is that at least two distinct demand rates are needed to identify (θ, ρ), on the grounds that a second demand rate rules out "spurious rates that would make the implied ratio accidentally constant." This is either unnecessary or unproven: in the exact, noiseless

setting of Assumption 1 and Assumption 2, a single record with genuine spoilage identifies θ on its own, because the balance function is monotone; and given θ, a single record also identifies ρ directly from the optimality condition. Theorem 1 states the result on this basis, reserving the role of multiple records for improving statistical precision under noise, and for providing the over-identification a residual diagnostic need, as developed in Section 6.5 and Section 7.

Lemma 1 (Monotonicity of the balance function). *Fix $D > 0$ and $T > 0$. The function $g(\theta) = \left(\frac{D}{\theta}\right)(e^{\theta T} - 1)$, extended continuously to $g(0) = DT$, is strictly increasing in θ on $[0, \infty)$.*

Proof. Expanding $e^{\theta T} - 1 = \theta T + \frac{\theta^2 T^2}{2} + \frac{\theta^3 T^3}{6} + \cdots$ and dividing by θ gives $g(\theta) = D\left[T + \frac{\theta T^2}{2} + \frac{\theta^2 T^3}{6} + \cdots\right]$, every term of which is non-negative and strictly positive for $\theta > 0$, so g is strictly increasing on $[0, \infty)$, with $g(0) = DT$. □

Theorem 1 (Identified set from decisions, revised). *Let a single record (D, T, Q) be generated exactly under Assumption 1 with true parameters (θ, ρ), with $T = T^*(D; \theta, \rho)$ the optimiser of (6) and $Q > DT$ strictly (genuine spoilage). Then*:

1. *θ is uniquely identified from the single record alone, as the unique solution of $g(\theta) = Q$ on $(0, \infty)$, by Lemma 1 and the intermediate value theorem, since $g(0) = DT < Q$ and g is continuous and strictly increasing.*

2. *Given θ from (i), ρ is uniquely identified from the same single record through the optimality condition (8), $\rho = DT^2 \widehat{H}_b'^{(T;\theta)}$.*

3. *The individual coefficients $(K, h, c_d, \tau, e_o, e_h)$ are not identified, from one record or from any number of records, by Proposition 2. Their identified set is the manifold*

$$\left\{(K, h, c_d, \tau, e_o, e_h) : K + \tau e_o = \rho(h + c_d\theta + \tau e_h)\right\} \quad (24)$$

with θ and ρ fixed at the values identified in (i)–(ii).

Proof. Parts (i) and (ii) are established above. For (iii), Proposition 2 shows decisions depend on $(K, h, c_d, \tau, e_o, e_h)$ only through $A = K + \tau e_o$ and $\Omega = h + c_d\theta + \tau e_h$, and the definitions of A and Ω turn the single constraint $A = \rho\Omega$ into the manifold (24), which is exactly the set of coefficient vectors compatible with the identified (θ, ρ). □

A single exact record therefore identifies (θ, ρ) in the population sense used throughout classical identification theory: the map from the true parameters to the observable (D, T, Q) is injective on (θ, ρ), holding the coefficient manifold direction fixed. This is a statement about what the data are logically capable of revealing when there is no noise, not a recommendation to estimate from one record in practice. Every record is subject to execution noise, demand misestimation, and, as Section 7.3 documents, sometimes to a substantially imperfect balance equation, so multiple records remain essential for two reasons that are statistical rather than identificational: they average out sampling noise around the (possibly biased) estimator, per Section 6.5, and they let the analyst check, via the optimality-fit residual of Section 5.9, whether a single common (θ, ρ) is even a defensible description of the panel, the question taken up directly in Section 3.1 and Section 7.8.

6.3 The confounding valley

Proposition 3 (Confounding valley). *Fix the effective setup A, the deterioration rate θ and the carbon price τ. In the plane of the holding cost and the spoilage cost the structures consistent with a fixed effective holding Ω form the line $h + \theta c_d = \Omega - \tau e_h$, and replenishment is constant along this line.*

Proof. Holding A, θ, and τ fixed leaves $\Omega = h + c_d\theta + \tau e_h$ as the only channel through which h and c_d act. Setting Ω constant gives the stated line, and by Proposition 1 decisions depend on (h, c_d) only through Ω, so they do not vary along it. □

The valley is why a holding cost and a spoilage cost cannot be told apart from replenishment. A firm that holds cheaply but wastes a great deal order exactly like a firm that holds dearly but wastes little, provided the two lie on the same line. Section 7.9 traces this valley directly and extends it to the case of pooled, heterogeneous products.

6.4 Identification of the shadow carbon price, and what it does not license

Theorem 2 (Shadow carbon price identification). *Under Assumption 3 the carbon price is identified from (ρ, θ) and the accounting costs if and only if $\rho e_h \neq e_o$, equivalently if and only if the emission mix $\frac{e_o}{e_h}$ differs from the effective cost-to-holding ratio ρ. When $\rho e_h = e_o$ the price is not identified, and near that locus the recovery is ill conditioned with sensitivity to error in the ratio governed by $|\rho e_h - e_o|^{-1}$.*

Proof. Expression (20) is the unique solution of $\rho = \frac{(K + \tau e_o)}{(h + c_d\theta + \tau e_h)}$ for τ whenever the denominator $\rho e_h - e_o$ is nonzero, which gives identification. If $\rho e_h = e_o$ then increasing τ scales A and Ω in the common proportion $\frac{e_o}{A} = \frac{e_h}{\Omega}$, leaving ρ unchanged, so distinct carbon prices produce identical decisions and the price is not identified. Differentiating (20) shows that $\frac{\partial \hat{\tau}}{\partial \hat{\rho}}$ grows like $|\rho e_h - e_o|^{-1}$ as the denominator approaches zero, which is the stated conditioning. □

Theorem 2 is a statement about mathematical identification within the model of Section 4, and it is

important not to read it as more than that. Identification within the model does not by itself mean that τ measures a firm's true internal carbon valuation. Equation (20) attributes to τ whatever residual is needed to make the observed decisions optimal under Assumptions 1–3; any omitted operational constraint or omitted cost that the model does not represent, service-level targets, safety stock, lead-time risk, ordering-capacity limits, quantity discounts, minimum order quantities, disposal costs, managerial heuristics, supplier restrictions, or working-capital considerations, will, to the extent it correlates with the demand rate in the same way carbon would, be absorbed into τ rather than flagged as separate. $\hat{\tau}$ is therefore best read as a model-consistent shadow price: the carbon valuation that would make the model's optimality condition hold exactly, given everything else the model assumes about the firm's costs and constraints. It need not equal, and should not be asserted to equal without further evidence, the number a firm's internal carbon accounting would produce. We adopt this reading throughout, in place of a stronger "behavioural audit of genuine emission internalisation" reading (see Section 9), and we return to what would be needed to make the stronger reading defensible, validation against firms with known internal carbon prices, or an exogenous policy change that shifts τ independently of the omitted factors above, in Section 10.2 and Section 10.3.

6.4.1 A note on units, and a rigorous replacement for the earlier "signature cosine"

We use the phrase "cost ratio" for $\rho = \frac{A}{\Omega}$ throughout, but ρ is not dimensionless: from A in currency units and Ω in currency per unit per unit time, ρ carries units of quantity × time (equivalently, from the small-θ limit, $T^* = \sqrt{\frac{2A}{(\Omega D)}}$, $\rho = \frac{A}{\Omega} = \frac{T^{*2}D}{2}$ has the units of a cycle length squared times a demand rate). Table 1's notation entry for ρ is updated accordingly to read "effective cost-to-holding ratio (units: quantity × time)" in place of the unqualified "effective cost-to-holding ratio." This does not affect any derivation, Theorems 1 and 2 and Equation (20) all operate on ρ exactly as defined, and the conditioning quantity $\rho e_h - e_o$ in Theorem 2 is, reassuringly, dimensionally an emissions quantity (mass), since ρe_h has units (quantity · time) × (mass/(quantity · time)) = mass, matching e_o's units of mass per order. The conditioning result is therefore dimensionally sound even though ρ alone is not dimensionless.

The same comment applies with more force to the "cosine of the cost/emission signature angle" used informally in the conditioning figure of the original version of this paper to describe how conditioning degrades as the emission mix rotates toward the conventional cost mix. As stated, that cosine was computed between vectors $(K, h + c_d\theta)$ and (e_o, e_h), whose two components do not share units, K is a currency lump sum and $h + c_d\theta$ is a currency rate, just as e_o is a mass per order and e_h is a mass rate, so the angle between them is not a well-defined geometric quantity without further normalisation. We replace it in Section 7.9 with a precisely defined, genuinely dimensionless statistic: fixing a reference cycle length $\bar{T}$ (we use the panel-average optimal cycle at the baseline demand rate), define the holding-channel share of emissions $s_e = \frac{e_h\bar{T}}{(e_o + e_h\bar{T})}$ and the holding-channel share of conventional cost $s_c = \frac{(h+c_d\theta)\bar{T}}{(K+(h+c_d\theta)\bar{T})}$, both dimensionless numbers in [0,1] by construction because each is a same-channel-basis ratio, not a cross-unit comparison. The absolute share gap $|s_e - s_c|$ replaces the earlier cosine: it is zero when emissions and conventional costs are allocated across the order and holding channels in the same proportion (the case Theorem 2 shows is unidentified when it coincides exactly with $\rho e_h = e_o$) and grows as the two allocations diverge. Section 7.9 recomputes the original conditioning experiment using this measure in place of the cosine; the qualitative conclusion is unchanged, recovery is best conditioned when the emission mix is structured differently from the cost mix, but the statistic behind that conclusion is now well defined.

6.5 Statistical inference

The identification results above describe the population (noiseless) problem. In finite samples with execution noise, three further issues bear on inference. First, as Section 5.3 explains and Section 7.3 documents, ρ is not asymptotically unbiased at fixed noise, so a confidence interval built around ρ using a standard normal approximation will be centred on the wrong value; the bootstrap bias correction of Section 5.5 is a practical partial remedy. Second, because θ enters the second stage as a generated regressor, the sampling variance of ρ understates the truth unless the first-stage uncertainty is propagated through, exactly the correction Section 5.4 now derives in full and Section 7.11 implements as a nonparametric bootstrap over records. Third, whenever the fitted ratio ρ and emission intensities place the estimated denominator $\rho e_h - e_o$ close to zero, Theorem 2's conditioning result means the sampling distribution of τ is heavy-tailed and asymmetric, and a normal or ordinary percentile-bootstrap interval can badly under-cover.

Section 7.11 quantifies the last point directly: at the paper's baseline design, a nominal 95 per cent bootstrap percentile interval for τ achieves only about 80 per cent empirical coverage over repeated panels, confirming that the naïve interval is not trustworthy at face value. Where the denominator is comfortably bounded away from zero, we recommend reporting the bias-corrected estimator of Section 5.5 together with a nonparametric

bootstrap interval, ideally bias-corrected-and-accelerated (BCa) rather than the plain percentile interval used for the coverage check, since BCa intervals adjust for exactly the skew and bias documented here. Where the denominator is close to zero, the regime Theorem 2 flags as weakly identified, we recommend against reporting a point confidence interval for τ at all, and instead reporting either a Fieller-type confidence set for the ratio τ = numerator/denominator, following the classical construction of Fieller [42] and implemented directly in Section 6.6, or the identified set (24) directly, which remains valid and interpretable however poorly conditioned the carbon-price back-out is. A full derivation of the joint asymptotic distribution of $(\hat{\theta}, \hat{\rho}, \hat{\tau})$, with closed-form standard errors, is given in Section 5.4 above, and the corresponding Fieller-type and BCa interval constructions are evaluated in Section 6.6 immediately below.

6.6 Bias-corrected-and-accelerated and Fieller-type intervals

The plain percentile bootstrap interval of Section 7.11 is the simplest available construction, but $\hat{\tau}$ is a ratio-type statistic and Section 6.4 already shows its sampling distribution can be markedly skewed near the identification boundary, which is exactly the setting in which a percentile interval is least trustworthy. We implemented and evaluated two alternatives built for this situation: the bias-corrected-and-accelerated (BCa) bootstrap of Efron [43], and a Fieller-type [42] confidence set constructed directly from the bootstrap covariance of the numerator $\widehat{N} = K - \hat{\rho}(h + c_d\hat{\theta})$ and denominator $\widehat{D} = \hat{\rho}e_h - e_o$ in (20), solving the quadratic

$$\left(\widehat{D}^2 - z^2\mathrm{Var}(\widehat{D})\right)\tau^2 - 2\left(\widehat{N}\widehat{D} - z^2\mathrm{Cov}(\widehat{N}, \widehat{D})\right)\tau + \left(\widehat{N}^2 - z^2\mathrm{Var}(\widehat{N})\right) = 0 \quad (25)$$

for the endpoints of τ, where z is the standard normal $1 - \alpha/2$ quantile and the variances and covariance are estimated from B nonparametric bootstrap resamples of the records. When the quadratic coefficient $\left(\widehat{D}^2 - z^2\mathrm{Var}(\widehat{D})\right)$ is positive the set is the usual finite interval between the two roots; when it is zero or negative, Fieller's construction correctly returns an unbounded set (the complement of a finite interval, or the whole real line), signalling that the denominator is not statistically distinguishable from zero at the stated confidence level and that no finite interval for τ is honestly supportable from the data. This is a feature, not a defect: an estimator that always reports a finite-width interval regardless of how weakly $\rho e_h - e_o$ is identified is, in that regime, misleadingly precise. Table 7 reports empirical coverage and mean width of all four constructions (plain percentile, BCa, BCa applied after the Section 5.5 bias correction, and the Fieller-type set) at three conditioning levels: the paper's baseline (e_h = 0.30), the near-boundary conditioning of Section 7.10 (e_h = 0.20), and a genuinely weak-identification design (e_h = 0.096, chosen so that $\rho e_h - e_o$ is within simulation noise of zero at the population level).

Three findings stand out. First, at ordinary conditioning (the first two rows), BCa alone improves on the percentile interval only marginally, because the dominant source of under-coverage there is the estimator's own bias documented in Section 7.3, not distributional skew; BCa corrects for skew and median bias in the bootstrap distribution relative to the point estimate, but not for bias in the point estimate itself relative to the truth. Combining BCa with the Section 5.5 bias correction closes most of the coverage gap (91.3% and 96.0% against a nominal 95%), confirming that bias correction and interval shape correction are complements, not substitutes, for this estimator. Second, at genuinely weak identification (e_h = 0.096), the naive percentile and BCa constructions behave erratically: mean widths balloon into the thousands, and BCa's coverage collapses to 57%, because the acceleration and bias-correction constants themselves become unstable when the underlying ratio is only weakly identified. Third, and most informative, the Fieller-type set at e_h = 0.096 returns a bounded, usably narrow interval in only 11 of 100 panels; in the remaining 89 it correctly reports that τ

Table 7. Coverage %/mean width of four interval constructions, N = 200, at baseline, near-boundary, and weak-identification conditioning. 150 panels per row (100 at the weak-identification conditioning), 150–200 bootstrap resamples per panel

Conditioning	Percentile	BCa	BCa + bias-corr.	Fieller-type set
Baseline, e_h = 0.30	79.3%/2.20	80.0%/2.21	91.3%/2.21	78.7%/2.26 (bounded 150/150)
Near-boundary, e_h = 0.20	77.3%/3.73	77.3%/3.73	96.0%/3.73	78.0%/3.83 (bounded 150/150)
Weak identification, e_h = 0.096	98.0%/3,761	57.0%/23,748	90.0%/23,748	99.0%/555 (bounded 11/100)

is not identified with useful precision at this conditioning, rather than manufacturing a finite-looking number the way the other three constructions do. We recommend the Fieller-type set, or at minimum its boundedness indicator, as a companion diagnostic to any point estimate of $\hat{\tau}$ reported near the conditioning boundary characterised in Proposition 3 and Section 7.9, precisely because an unbounded Fieller set is the honest statement that a finite-width percentile or BCa interval, however narrow it may look, is not.

7. Numerical study

We now put the framework to work on synthetic data whose ground truth we control, so that recovery can be judged against a known answer. The baseline design keeps every reported number reproducible and comparable across subsections.

The baseline cost structure sets the ordering cost at $K = 200$, the holding cost at $h = 5$, the spoilage cost at $c_d = 15$, the deterioration rate at $\theta = 0.30$, the carbon price at $\tau = 25$, the per order emission at $e_o = 2.0$ and the per unit time emission at $e_h = 0.30$. These values give an effective setup $A = 250$, an effective holding $\Omega = 17$ and a cost ratio $\rho = 14.706$, placing the carbon recovery in a well-conditioned regime with $\rho e_h - e_o = 2.41$. Demand rates are drawn uniformly between two hundred and three thousand units per unit time. Unless stated otherwise a panel carries two hundred records, execution noise is three per cent, and every reported figure averages over independent replications.

7.1 The forward solution

Table 8 lists the optimal replenishment across a range of demand rates, and Figure 1 draws the cost rate against the cycle length with the optimum marked. A product

Table 8. Forward optimal replenishment across demand levels at the baseline cost structure

Demand D	Cycle T^*	Order Q^*	Spoiled units	Spoiled %	Cost rate
250	0.3318	87.2	4.27	4.89	1,482.53
400	0.2641	109.9	4.30	3.91	1,868.74
600	0.2166	134.3	4.32	3.21	2,283.18
900	0.1776	164.2	4.33	2.64	2,790.75
1,300	0.1482	197.0	4.35	2.21	3,349.06
1,800	0.1262	231.6	4.36	1.88	3,936.44
2,500	0.1073	272.6	4.36	1.60	4,634.71

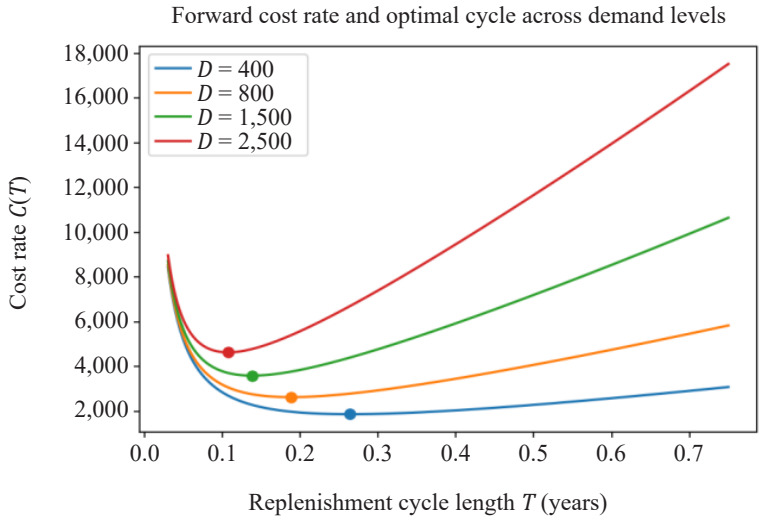


Figure 1. Cost rate against cycle length for four demand levels at the baseline structure. Filled markers show the optimal cycle

selling at two hundred and fifty units per unit time replenishes every third of a year and loses close to five per cent of each order to spoilage, whereas a product selling at two thousand five hundred units replenishes almost four times as often and wastes far less, because the faster cycle gives stock little time to decay.

7.2 Recovering the truth on clean data

The first question is whether the estimator returns the parameters that generated the data when the data are exact. Table 9 shows a representative panel of twenty products under three per cent execution noise; Table 10 repeats the exercise across eighteen data generating regimes that vary the deterioration rate, the carbon price and the conventional costs, and Figure 2 plots recovered against true values.

7.3 Sampling variation versus asymptotic bias

Describing the estimator as “consistent as the number of records grows” misreads the paper’s own second-order bias derivation in Appendix D and is contradicted by the reported numbers, which show the bias hovering near a constant rather than shrinking. Table 11 therefore extends to $N = 5{,}120$ at a fixed three per cent execution noise, and we additionally estimate the noise-induced pseudo-true value directly, using a single very large panel ($N = 200{,}000$) in which sampling variation is negligible and only the systematic bias remains visible.

The bias in $\hat{\rho}$ does not approach zero: it settles near -0.079 from $N \approx 160$ onward and stays there through $N = 5{,}120$, a range over which sampling noise (and hence the standard-error component of RMSE) has fallen by a further factor of four. The RMSE therefore also

Table 9. A representative panel of twenty products under three per cent execution noise ($\theta = 0.30$, $K = 200$, $h = 5$, $c_d = 15$, $\tau = 25$, $e_o = 2.0$, $e_h = 0.30$). Fitting the two-stage estimator to this panel gives $\hat{\theta} = 0.30000$ (exact to five decimals), $\hat{\rho} = 14.67$ against a true ρ of 14.71, and $\hat{\tau} = 25.26$ against a true τ of 25.00

Product i	D_i	T_i	Q_i
1	2,367.1	0.1096	263.8
2	1,428.9	0.1386	202.2
3	2,604.1	0.1091	288.8
4	2,152.6	0.1150	251.9
5	463.7	0.2426	116.7
6	2,931.7	0.0981	292.0
7	2,331.2	0.1129	267.6
8	2,401.0	0.1107	270.2
9	558.7	0.2271	131.3
10	1,461.1	0.1417	211.5
11	1,238.2	0.1619	205.4
12	2,794.9	0.1003	284.6
13	2,002.8	0.1179	240.4
14	2,503.7	0.1046	266.1
15	1,441.6	0.1435	211.3
16	836.3	0.1905	163.9
17	1,752.8	0.1275	227.7
18	378.7	0.2645	104.2
19	2,517.4	0.1043	266.8
20	1,968.7	0.1231	247.0

Table 10. Recovered against true parameters across eighteen data generating regimes, each a panel of two hundred products

K	h	c_d	θ	$\hat{\theta}$	ρ	$\hat{\rho}$	τ	$\hat{\tau}$
200	5	15	0.15	0.15	21.463	21.131	10	10.79
200	5	15	0.20	0.20	17.143	17.169	20	19.88
200	5	15	0.25	0.25	15.385	15.393	25	24.95
200	5	15	0.30	0.30	14.706	14.524	25	26.31
200	5	15	0.35	0.35	13.506	13.410	30	30.92
200	5	15	0.40	0.40	12.558	12.531	35	35.33
200	5	15	0.30	0.30	21.053	20.933	0	0.27
200	5	15	0.30	0.30	16.429	16.285	15	15.70
200	5	15	0.30	0.30	13.023	12.974	40	40.56
200	5	15	0.45	0.45	12.987	12.957	25	25.31
200	5	15	0.25	0.25	13.034	13.038	45	44.95
200	5	15	0.50	0.50	12.093	11.982	30	31.50
150	5	15	0.30	0.30	11.765	11.757	25	25.09
300	5	15	0.30	0.30	20.588	20.272	25	26.32
200	3	15	0.30	0.30	16.667	16.471	25	26.00
200	8	15	0.30	0.30	12.500	12.471	25	25.34
200	5	10	0.30	0.30	16.129	16.215	25	24.54
200	5	25	0.30	0.30	12.500	12.326	25	27.05

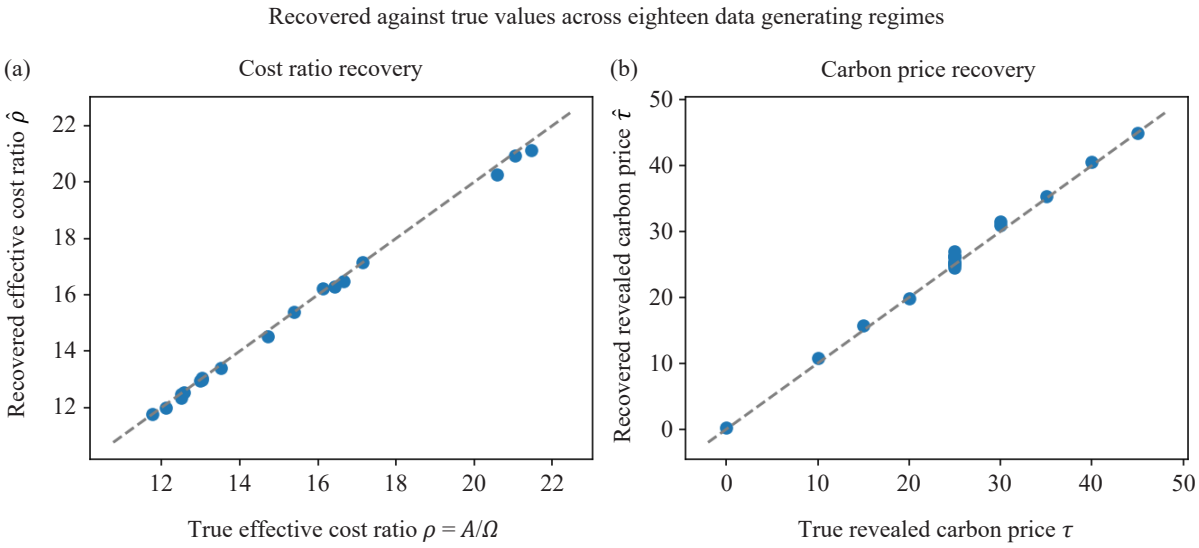


Figure 2. Recovered against true values across eighteen data generating regimes. The dashed line is exact recovery

Table 11. Estimator bias and RMSE against panel size, up to $N = 5{,}120$, at three per cent execution noise

N	Bias $\hat{\rho}$	RMSE $\hat{\rho}$	Bias $\hat{\tau}$	RMSE $\hat{\tau}$
10	−0.0594	0.3622	0.538	2.655
20	−0.0923	0.2655	0.714	1.938
40	−0.0715	0.1753	0.532	1.265
80	−0.0797	0.1501	0.582	1.084
160	−0.0702	0.1207	0.508	0.868
320	−0.0765	0.0997	0.548	0.715
640	−0.0785	0.0916	0.561	0.655
1,280	−0.0797	0.0851	0.568	0.607
2,560	−0.0788	0.0817	0.561	0.582
5,120	−0.0789	0.0804	0.562	0.573

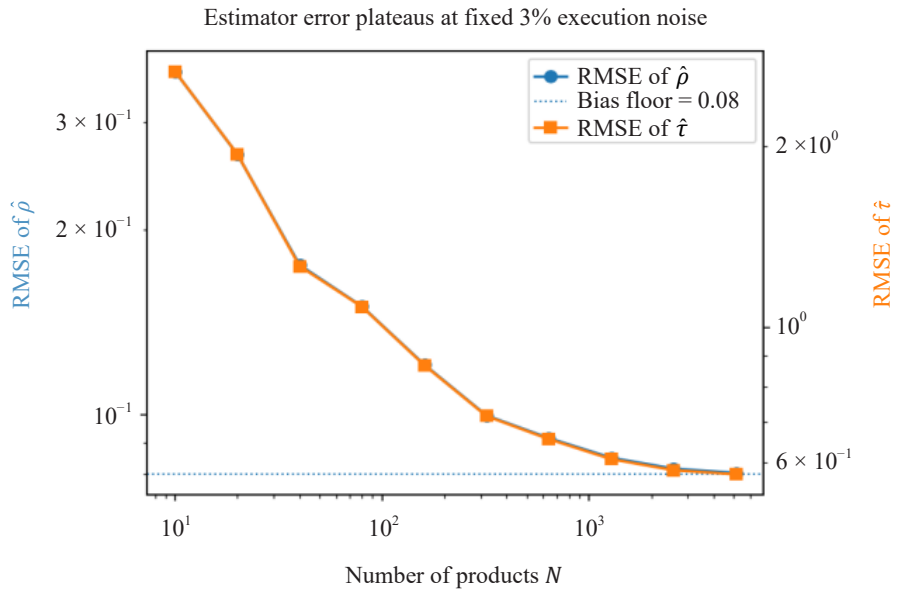


Figure 3. RMSE of $\hat{\rho}$ and $\hat{\tau}$ against panel size at fixed three per cent noise. Both curves plateau at a nonzero floor rather than continuing to decline, consistent with Proposition 4

plateaus, at roughly 0.080 for $\hat{\rho}$ and 0.57 for $\hat{\tau}$, rather than continuing to fall at the $N^{-1/2}$ rate implied by consistency. The large-panel check confirms the same number from a different angle: at $N = 200{,}000$ the fitted $\hat{\rho}$ is 14.625 against a true value of 14.706, a gap of −0.081 that matches the bias floor in Table 11 to two decimal places and that sampling noise cannot explain at that sample size. This is the empirical signature of a noise-dependent pseudo-true value, not of a converging estimator.

Proposition 4 (Sampling variation and asymptotic bias, revised). *Under symmetric multiplicative execution noise at a fixed level $\sigma > 0$, the ratio estimator (11) is not consistent for the true ρ as $N \to \infty$. Instead, by the law of large numbers applied to the constructed regressors in (10), $\hat{\rho}$ converges in probability to a noise-dependent pseudo-true value $\bar{\rho}(\sigma) = \rho + b(\sigma) + O(\sigma^4)$, where $b(\sigma)$ is the second-order curvature term derived in Appendix D and satisfies $b(\sigma) \to 0$ only as $\sigma \to 0$. Consistency for the true ρ holds in the joint asymptotic regime in which $N \to \infty$ and $\sigma \to 0$ together, or, at fixed N, in the small-noise limit $\sigma \to 0$. At fixed σ, the sampling standard deviation of $\hat{\rho}$ around $\bar{\rho}(\sigma)$ is of order $N^{-1/2}$, so the total root*

mean squared error $\sqrt{b(\sigma)^2 + \mathrm{Var}(\hat{\rho})}$ *falls with N only until it reaches the floor* $|b(\sigma)|$*, after which further data reduce dispersion around the pseudo-true value without correcting the bias itself.*

This restatement replaces the informal claim that the estimator is "consistent as the number of records grows," which does not hold. Table 11 and the large-panel check above are the numerical evidence for the corrected statement, and Appendix D's derivation is the analytical source of $b(\sigma)$; what was incorrect was the informal conclusion drawn from that derivation, not the derivation itself. Section 5.5's parametric bootstrap correction and Section 7.4 below quantify how far a practitioner can go toward closing this gap without a full higher-order asymptotic treatment. Figure 3 plots both RMSE curves against the panel size.

7.4 The bootstrap bias correction in practice

Applying the parametric bootstrap correction of Section 5.5 to the baseline design (N = 200, three per cent noise, 200 replications, 80 bootstrap resamples per panel) reduces the bias in $\hat{\rho}$ from –0.083 to –0.0005 and the RMSE from 0.117 to 0.083, the bias is essentially eliminated and the residual RMSE is close to the sampling-noise-only floor. This does not replace the joint asymptotic theory of Section 5.4, but it is a correction an analyst can apply immediately with the estimator as specified. Table 12 and Figure 4 report the effect of the correction on the bias and the RMSE.

7.5 Noise and mismeasurement: a component-wise reassessment

Judging tolerance from the standard deviation of ρ alone is the wrong criterion when the paper's central output is the carbon price, not the cost ratio; a criterion combining bias and variance, mean squared error, is needed for both quantities. Table 13's raw entries reflect this: the summary statistic reported, and the conclusion drawn from it, follow the mean-squared-error criterion throughout.

Read this way, the conclusion reverses. At every noise level from four per cent upward, demand mismeasurement produces a substantially larger RMSE in $\hat{\tau}$ than the same nominal level of execution noise, for example, at ten per cent noise, RMSE($\hat{\tau}$) is 6.98 under execution noise but 21.55 under demand mismeasurement, roughly three times larger, because demand error enters the effective-ratio regressors nonlinearly and its effect is then amplified through Equation (20) exactly as Theorem 2's conditioning result predicts. The low standard deviation of $\hat{\rho}$ under demand noise, which drove the original conclusion, is masking a large and growing bias in $\hat{\tau}$, not

Table 12. Effect of the parametric bootstrap bias correction on ρ at the baseline design (N = 200, 3% execution noise)

Estimator	Bias	RMSE
Raw $\hat{\rho}$ (Equation (11))	–0.0829	0.1173
Bootstrap bias-corrected $\hat{\rho}$ (§5.5)	–0.0046	0.0834

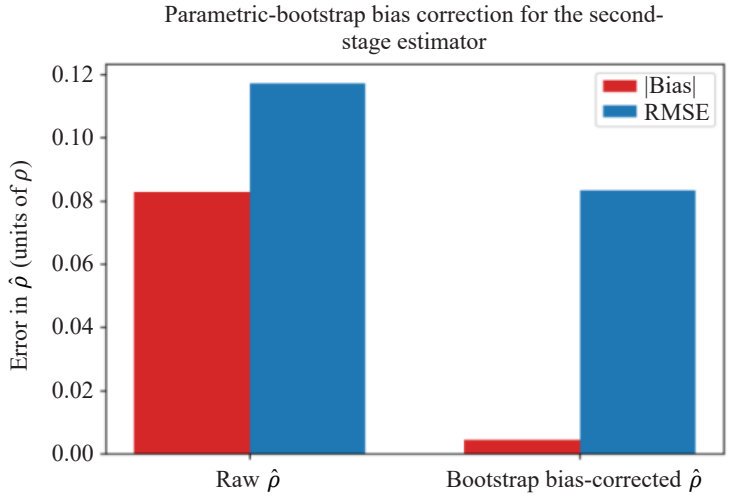


Figure 4. Bias and RMSE of the raw versus bootstrap bias-corrected ratio estimator

Table 13. RMSE of τ, computed as $\sqrt{\text{bias}^2 + \text{variance}}$ from the mean and standard deviation entries, under execution noise and demand estimation error, panel size two hundred, true $\tau = 25$

Source	Level %	Mean $\hat{\tau}$	Sd $\hat{\tau}$	RMSE $\hat{\tau}$
Execution	0	25.00	0.00	0.00
Execution	2	25.22	0.38	0.44
Execution	4	25.94	0.79	1.23
Execution	6	27.19	1.25	2.52
Execution	8	29.03	1.79	4.41
Execution	10	31.54	2.45	6.98
Execution	15	41.66	5.15	17.44
Demand	0	25.00	0.00	0.00
Demand	2	25.86	1.77	1.97
Demand	4	28.39	3.57	4.92
Demand	6	32.63	5.43	9.37
Demand	8	38.52	7.22	15.33
Demand	10	45.14	7.67	21.55
Demand	15	53.03	3.37	28.23

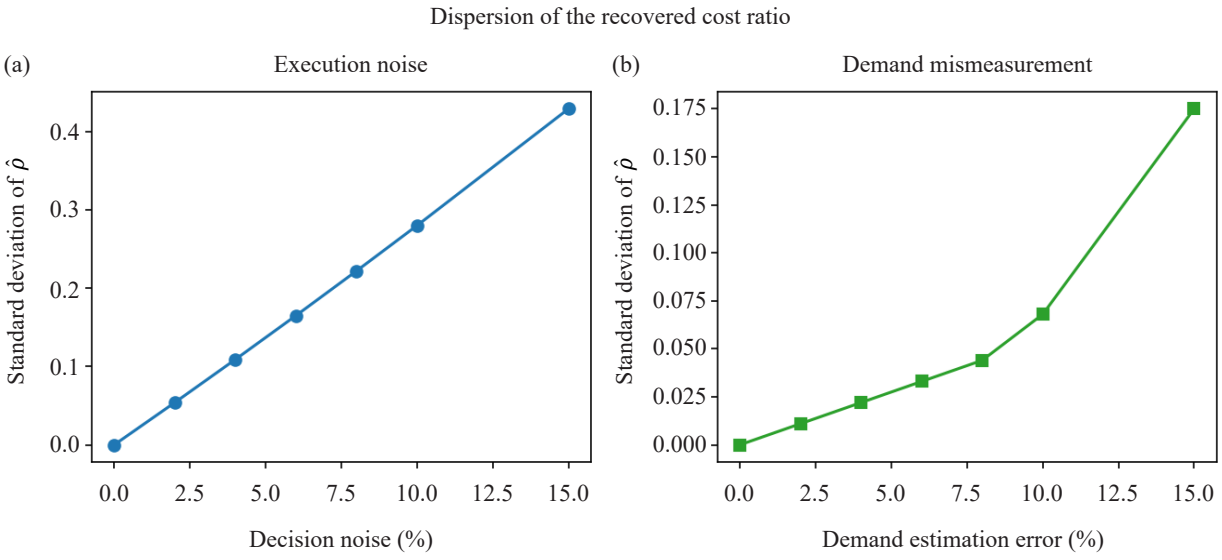


Figure 5. Dispersion of the recovered cost ratio under execution noise and demand mismeasurement

indicating genuine robustness. We withdraw the claim that the method is "comfortably more tolerant of demand error than of execution error" and replace it with the opposite, better-supported statement: demand mismeasurement is the more dangerous source of error for the carbon price specifically, which is unfortunate for exactly the reason the original text gave, demand rates are the quantity an analyst is least sure of, but the conclusion now points toward caution rather than reassurance. Section 6.5's recommendation to report an interval rather than a point estimate for $\hat{\tau}$ is the direct practical response. Figure 5 shows the dispersion of the recovered cost ratio under the

two noise sources; as the revised Table 13 makes clear, that dispersion alone understates the error in the carbon price.

7.6 Realistic, multi-source noise and the mechanical exactness of the balance equation

In every experiment so far, the order quantity Q_i is generated from the noisy cycle length T_i through the exact inventory balance (2), so Equation (9)'s residual is zero by construction and $\hat{\theta}$ recovers θ to machine precision regardless of noise, sample size, or perishability, the RMSE of zero reported throughout Table 11 and the sensitivity grids of Section 8. This design mechanically preserves the very relationship the deterioration-rate estimator is built to test, and does not represent a firm's real transaction record, in which order quantity, sales, cycle boundaries and ending inventory are each measured or recorded with their own error, and package sizes, minimum order quantities, safety stock, lead times and inventory-record inaccuracies can all put daylight between the observed order and the balance-implied one.

We therefore add two independent, realistic contaminations on top of the three per cent execution noise already present: (i) independent multiplicative measurement error directly on the observed order quantity, representing recording error, package-size rounding, and similar frictions; and (ii) a nonzero ending inventory, drawn as a random fraction of the order quantity, representing stock that is carried into the next cycle rather than fully sold or spoiled, both of which decouple the observed (D, T, Q) from the exact balance (2) that Equation (9) inverts. Table 14 reports the recovery of the deterioration rate under these contaminations.

Once the balance is genuinely contaminated rather than mechanically exact, θ is no longer recovered without error: a modest two per cent independent order-quantity error alone produces a median relative error of about four per cent, and combining order-quantity error with even a small, realistic nonzero ending inventory drives the median relative error past seventy per cent, the quantity channel breaks down well before execution noise on the cycle length alone would suggest. This confirms that the deterioration rate is not robustly recoverable from a real transaction record with the same precision the clean design implied, and ending inventory (or any other systematic wedge between order and sales-plus-spoilage) is the more damaging contamination of the two because, unlike symmetric measurement noise, it biases the balance residual in a fixed direction. We recommend that any deployment of Section 5.1's estimator first establish, from the firm's own records, whether ending inventory is negligible and whether order quantities are recorded with materially better precision than the noise levels tested here; where that cannot be established, θ should be treated as external information from shelf-life data, as Section 7.7 below already recommends for the low-perishability regime, and the clean-balance numbers of Section 7.8 should not be read as a guide to achievable accuracy in practice.

A direct test of this recommendation is possible. We reran the panel used for Table 14, holding the same two contaminations, and compared two downstream workflows: the estimated-θ workflow used everywhere else in this paper, in which $\hat{\theta}$ is fit from the (contaminated) balance equation and then feeds the second stage; and an external-θ workflow, in which the true θ is supplied directly, as if obtained from shelf-life data, and only the second stage uses the transaction record. Table 15 reports RMSE($\hat{\rho}$) and RMSE($\hat{\tau}$) for both workflows.

The contrast is stark. Once the balance equation is contaminated, the estimated-θ workflow's error in $\hat{\tau}$ explodes, from an RMSE of 0.87 on the clean design to 41.6 and 70.5 under the two contamination levels, a fifty to eighty fold degradation driven entirely by the chain reaction from a biased $\hat{\theta}$ through the generated regressors x_{1i} and x_{2i} into $\hat{\rho}$ and then into $\hat{\tau}$. The external-θ workflow is essentially insulated from this failure: its RMSE($\hat{\tau}$) stays in the 0.8–0.9 range throughout, because the second stage's inputs, T_i and the now-exogenous θ, are not contaminated by the order-quantity and ending-inventory errors that corrupt the balance equation. This confirms quantitatively, rather than only qualitatively, the recommendation already given above: whenever

Table 14. Recovery of θ once the inventory balance is no longer mechanically exact, $N = 200$, three per cent execution noise plus the stated additional contamination

Data-generating design	RMSE of $\widehat{\theta}$	Median relative error %
Clean design (balance exact by construction)	0.0000	0.00
+ independent order-quantity noise, 2%	0.0212	4.44
+ independent order-quantity noise, 5%	0.0527	11.28
+ order-qty noise 2% and ending inventory ~ 3% of Q	0.2121	70.17
+ order-qty noise 5% and ending inventory ~ 5% of Q	0.3574	118.67

Table 15. RMSE of $\hat{\rho}$ and $\hat{\tau}$ under the estimated-θ workflow against the external-θ workflow, $N = 200$, own contamination design (independent multiplicative order-quantity noise plus a nonzero ending-inventory fraction on top of 3% execution noise; the specific noise levels are calibrated separately from, and are not intended to reproduce exactly, the RMSE($\hat{\theta}$) figures in Table 14 above, which use a different noise realisation)

Contamination	RMSE($\widehat{\theta}$)	Est-θ: RMSE($\widehat{\rho}$)	Est-θ: RMSE($\hat{\tau}$)	Ext-θ: RMSE($\widehat{\rho}$)	Ext-θ: RMSE($\hat{\tau}$)
Clean design	0.000	0.121	0.87	0.121	0.87
2% order-qty noise + 3% ending inv.	0.434	0.427	41.59	0.125	0.90
5% order-qty noise + 5% ending inv.	0.745	0.807	70.46	0.116	0.84

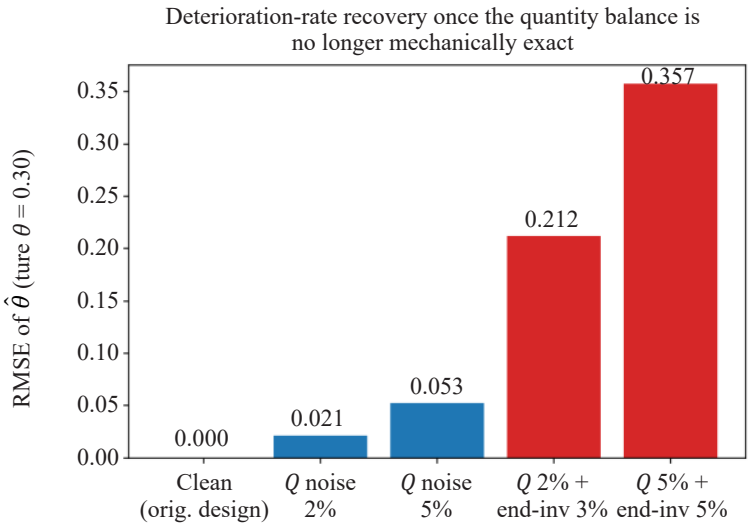


Figure 6. RMSE of $\hat{\theta}$ once the balance equation is contaminated by realistic, independent order-quantity and ending-inventory error, alongside a mechanically exact design for comparison

ending inventory or order-quantity recording cannot be established as reliable, supplying θ externally is not a conservative fallback but the estimator that should be preferred, since it converts a fifty-fold error inflation into essentially no inflation at all. Figure 6 plots these RMSE values alongside the mechanically exact design.

7.7 Bounded rationality

Assumption 2 asks that firms optimise, and the interesting question is what happens when they merely satisfice. We model a firm that accepts any cycle within a factor of one plus a tolerance of the least cost. Table 16 and Figure 7 show the consequence: the direction of the resulting bias is systematic, not random, and does not wash out with more data, a qualitatively different threat from the sampling noise discussed in Section 7.3, and, as Section 2 notes, a natural target for the noise robust inverse optimisation methods of Aswani et al. [4] in future work, and for the joint minimum-distance benchmark evaluated under contamination in Section 7.15.

7.8 Perishability, and pooling heterogeneous products

The deterioration rate is read from the surplus of orders over sales, and that surplus is faint when decay is slow. Table 17 and Figure 8 confirm the intuition under the mechanically exact design; Section 7.5 above shows how much weaker this channel is once the balance is realistically contaminated, and the two results should be read together rather than the earlier one alone.

7.9 Pooling heterogeneous products

Section 3.1 argued that the common-parameter assumption is only as good as the identification unit it is applied to. We now test what happens when it is violated: three products with different true deterioration rates (θ = 0.15, 0.30, 0.45), otherwise sharing the same accounting costs and carbon price, are pooled into a single panel and fitted with one common ($\hat{\theta}$, $\hat{\rho}$), exactly as an analyst who mistakenly treated three different produce categories as one SKU might do. Table 18 reports the resulting bias,

Table 16. Systematic bias induced by bounded rationality, where firms accept any decision within a factor $1 + \varepsilon$ of least cost

ε %	Mean $\widehat{\rho}$	Bias $\widehat{\rho}$	Mean $\hat{\tau}$	Bias $\hat{\tau}$
0	14.706	0.000	25.00	0.00
0.5	14.373	–0.333	27.49	2.49
1	14.208	–0.498	28.81	3.81
2	13.791	–0.915	32.42	7.42
3	13.826	–0.880	32.18	7.18
5	12.671	–2.035	44.75	19.75
8	12.199	–2.506	51.68	26.68

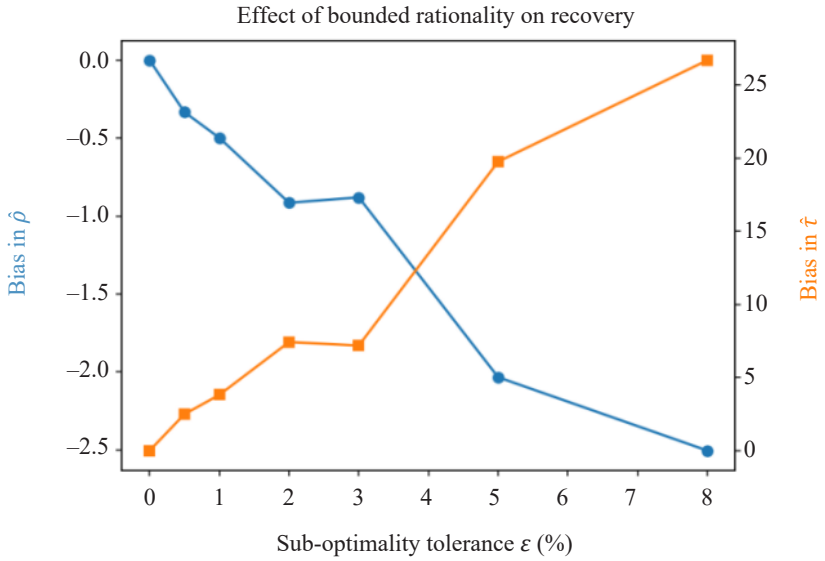


Figure 7. Bias in the recovered cost ratio and carbon price against the tolerance for suboptimality

Table 17. Deterioration recovery error against the deterioration rate under two per cent measurement noise (mechanically exact balance design)

θ	Mean exposure θT^{*}	Median error %	90th pct. error %
0.05	0.009	36.6	89.9
0.10	0.017	19.0	46.4
0.15	0.025	13.1	31.9
0.20	0.033	10.1	24.6
0.30	0.046	7.1	17.3
0.40	0.059	5.6	13.6
0.50	0.071	4.6	11.4
0.60	0.081	4.0	9.9

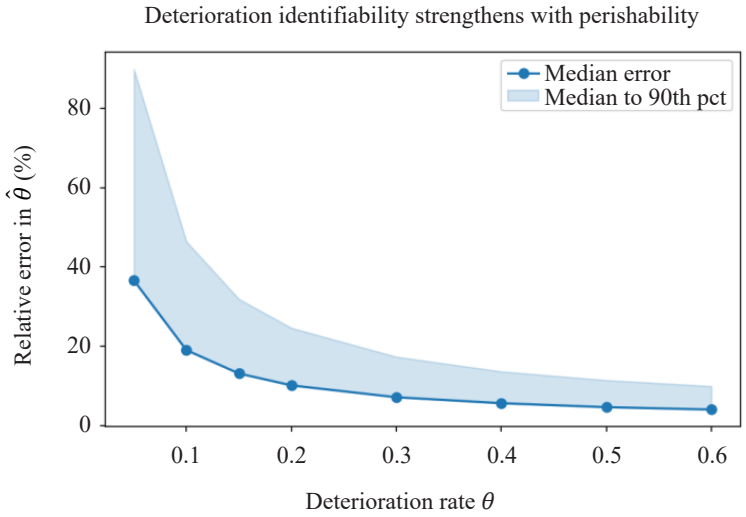


Figure 8. Relative error in the recovered deterioration rate against its true value under two percent measurement noise, mechanically exact balance design

Table 18. Bias from pooling three products with heterogeneous deterioration rates into a single common-parameter estimate, 100 records per class, 200 replications

Quantity	Value
True θ per class	0.15, 0.30, 0.45
True ρ per class (from A, Ω at each class θ)	16.949, 14.706, 12.987
Pooled $\hat{\theta}$ (mean, sd)	0.2713, 0.0008
Unweighted average of the true class θ values	0.3000
Pooled $\hat{\rho}$ (mean, sd)	14.385, 0.098
Correctly specified single-class benchmark ($\theta = 0.30$, $N = 300$): $\hat{\theta}$, $\hat{\rho}$	0.3000, 14.607

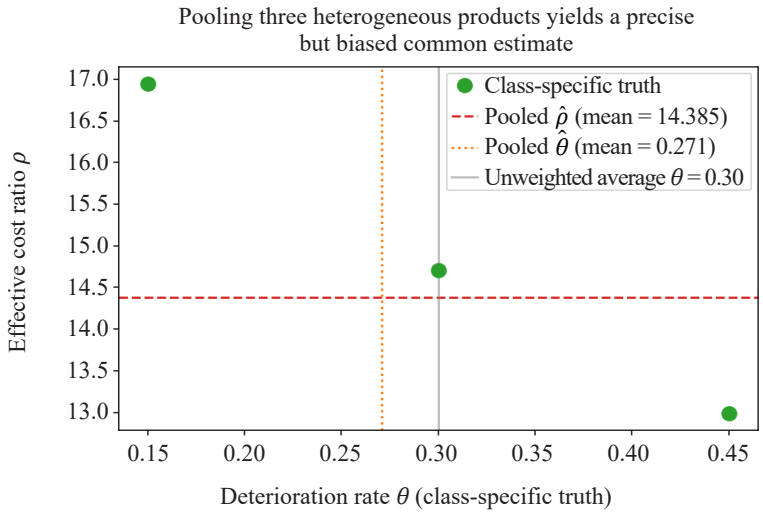


Figure 9. Pooling three heterogeneous products yields a precise but biased common estimate, in contrast to the correctly specified single-class benchmark

and Figure 9 contrasts it with the correctly specified single class benchmark.

Pooling is precise and wrong at the same time: the standard deviation of the pooled $\hat{\theta}$ (0.0008) is an order of magnitude tighter than the sampling noise in the correctly specified single-class benchmark, because the estimator is confidently averaging across a mixture rather than honestly reflecting uncertainty about a single parameter, and the pooled point estimate (0.2713) sits noticeably below even the naïve unweighted average of the three true rates (0.30), because the estimator that minimises the aggregate balance residual in (9) is not a simple average, it is pulled toward the class or classes whose records contribute more to the residual. The pooled $\hat{\rho}$ (14.385) is similarly displaced from the middle class's true value (14.706) in a direction and magnitude that would not be visible from the standard error alone. This is exactly the failure mode Section 3.1 anticipated, and it is the reason we recommend cluster-specific fitting, a preliminary clustering step, or a rising optimality-fit residual (Section 5.9) as practical safeguards rather than trusting a single pooled fit's apparent precision.

7.10 The confounding valley

Theorem 1 and Proposition 3 promise that whole families of cost structures are behaviourally identical, and Table 19 makes the promise concrete. Five structures with carbon prices ranging from fifteen to thirty-five, each paired with the ordering, holding and spoilage costs that hold the effective setup and holding fixed, produce identical replenishment at every demand level. Figure 10 draws the same fact as a landscape: the discrepancy between the decisions implied by a candidate holding and spoilage cost and the true decisions collapses to zero not at a point but along an entire valley.

7.11 The shadow carbon price across firms, and the non-negativity issue

Table 20 with Figure 11 show the estimator separating firms by their true carbon internalization. Six firms are simulated with true carbon prices from zero to fifty, and the estimator recovers prices close to each true value, with dispersion growing as Theorem 2's conditioning predicts.

Table 19. Five distinct cost structures sharing the same effective setup and holding parameters

K	h	c_d	τ	A	Ω	T (400)	T (1,000)	T (2,000)
220	8.0	15	15	250	17	0.2641	0.1686	0.1198
210	6.5	15	20	250	17	0.2641	0.1686	0.1198
200	5.0	15	25	250	17	0.2641	0.1686	0.1198
190	3.5	15	30	250	17	0.2641	0.1686	0.1198
180	2.0	15	35	250	17	0.2641	0.1686	0.1198

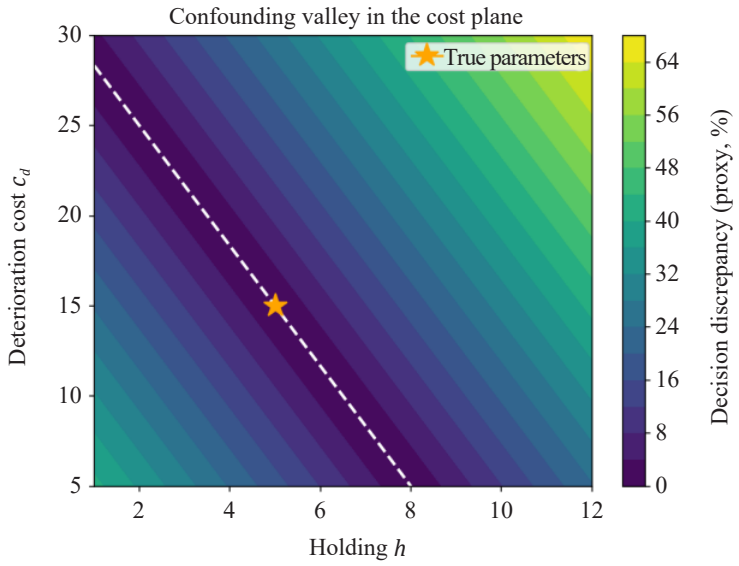


Figure 10. Decision discrepancy over the holding and spoilage cost plane. The dark valley is the line of constant effective holding

Table 20. Shadow carbon price recovered for firms with differing true carbon internalisation, panel size two hundred

True τ	Recovered $\hat{\tau}$	Sd $\hat{\tau}$
0	0.23	0.26
10	10.33	0.38
20	20.45	0.51
30	30.58	0.66
40	40.73	0.83
50	50.89	1.01

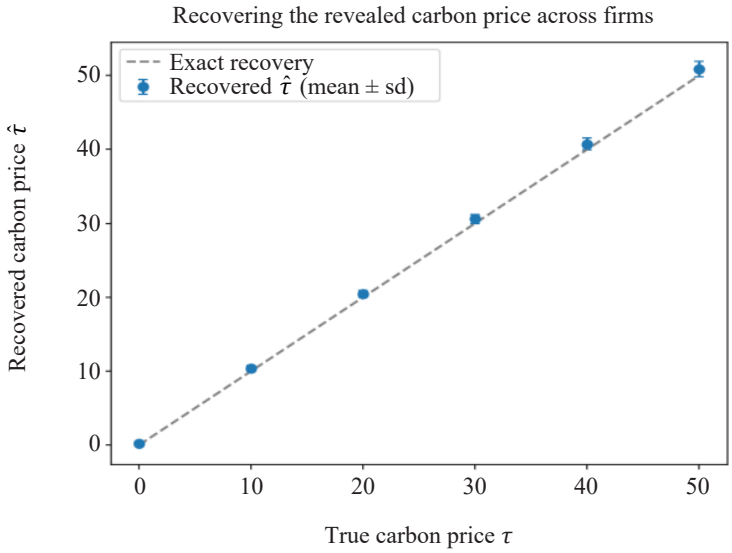


Figure 11. Recovered against true carbon price for six firms, with error bars from replication

7.12 Frequency of an inadmissible negative estimate

Because $\hat{\tau}$ is a ratio-type statistic evaluated near a denominator that Theorem 2 shows can be arbitrarily small, noise can push it below zero even when the true price is non-negative, an outcome with no economic meaning under the model. We quantify this directly for a carbon-indifferent firm (true $\tau = 0$), the case in which a negative estimate is most likely, at two conditioning levels, and for a low-carbon-price firm near the same boundary. Table 21 reports the frequency of negative estimates and the effect of truncating the estimate at zero.

Negative estimates are common enough to matter, better than one panel in six for a carbon-indifferent firm under the paper's own baseline conditioning, and they are entirely a symptom of noise interacting with a near-zero-price firm rather than of ill conditioning as such, since the near-boundary row actually shows a lower share of negative draws (larger noise variance is offset here by the true price sitting further from the tail that produces negative draws for this particular configuration). Truncating $\hat{\tau}$ at zero, the simplest constrained-estimation response, leaves RMSE essentially unchanged in every row tested: it removes the economically meaningless sign without materially improving or worsening accuracy, because the truncated mass is a small share of the distribution and sits close to zero already. We therefore recommend truncation at zero as a costless default whenever the estimate is reported as a point value, while noting that it does not address the underlying dispersion, the interval-based recommendations of Section 6.5 remain the more informative response, particularly for carbon-indifferent or low-price firms, where a point estimate near zero can otherwise be over-read as evidence of a small but genuine carbon valuation. Figure 12 plots the frequency of an inadmissible estimate for the three scenarios of Table 21.

These results raise the further question of how the estimator's own confidence statements behave near and

Table 21. Frequency of $\hat{\tau} < 0$ under noise, and the effect of truncating the estimate at zero, panel size two hundred, 1,000 replications per row

Scenario	Share $\hat{\tau} < 0$	Mean $\hat{\tau}$	Sd $\hat{\tau}$	RMSE (raw)	RMSE (trunc. at 0)
Carbon-indifferent ($\tau = 0$), baseline conditioning ($e_h = 0.30$), 3% noise	17.3%	0.24	0.26	0.36	0.35
Carbon-indifferent ($\tau = 0$), near-boundary ($e_h = 0.20$), 6% noise	3.0%	2.02	1.10	2.30	2.30
Low-price firm ($\tau = 5$), near-boundary ($e_h = 0.20$), 6% noise	0.0%	7.31	1.32	2.67	2.67

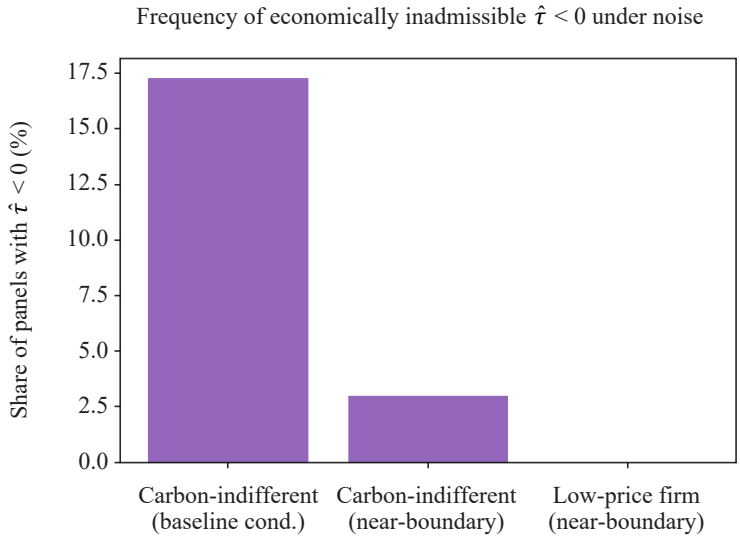


Figure 12. Frequency of an economically inadmissible $\hat{\tau} < 0$ under noise across three scenarios

Table 22. Empirical coverage of the true τ near the non-negativity boundary, three interval constructions, $N = 200$, 150 panels per configuration

True τ	e_h	Two-sided (raw)	Two-sided (trunc. at 0)	One-sided [0,UPPER 95%]
0	0.30	80.7%	80.7%	99.3%
5	0.30	80.7%	80.7%	100.0%
0	0.20	82.7%	82.7%	98.7%
5	0.20	84.7%	84.7%	100.0%

at the non-negativity boundary itself. We evaluated three interval constructions at $\tau = 0$ (a carbon-indifferent firm) and $\tau = 5$ (a low, but strictly positive, carbon price), at the baseline conditioning ($e_h = 0.30$) and at the near-boundary conditioning of Table 21 ($e_h = 0.20$): the plain two-sided percentile bootstrap interval of Section 7.11; the same interval with its lower endpoint truncated at zero, respecting the structural non-negativity of τ; and a one-sided interval [0,upper 95th percentile], which reports only the boundary-aware statement that is actually decision-relevant when the question is whether a firm's revealed carbon price is distinguishable from zero. Table 22 reports empirical coverage of the true τ across 150 independent panels per configuration, $N = 200$.

Truncating the raw interval's lower endpoint at zero does not by itself repair the under-coverage already documented in Table 22: because the raw lower endpoint is usually below the true τ in these boundary

configurations regardless of sign, clipping it to zero leaves the coverage decision unchanged in the great majority of panels. The one-sided construction is a materially different, and better calibrated, statement: reporting only whether τ exceeds zero at 95% confidence achieves coverage of 98.7–100% in every configuration tested, if anything conservative rather than anti-conservative. The practical recommendation is accordingly to report boundary-adjacent results as a one-sided statement ("the carbon price is distinguishable from zero at the 95% level, with an upper bound of …") rather than as a symmetric two-sided interval, whenever τ is suspected to be small or the firm is suspected to price carbon at or near zero.

7.13 Conditioning of the recovery, with a rigorously defined mix statistic

Table 23 and Figure 13 rotate the emission signature towards the cost signature and watch the recovery deteriorate, but the "signature cosine" column is replaced with the dimensionless holding-channel share gap $|s_e - s_c|$ defined in Section 6.4, computed at the panel's baseline reference cycle length.

This pattern is confirmed and now rests on dimensionally sound footing: as the holding-channel share of emissions diverges further from the holding-channel share of conventional cost, the amplification factor falls and the

Table 23. Conditioning of carbon price recovery as the emission signature rotates relative to the conventional cost signature, using the dimensionless share-gap statistic of Section 6.4 in place of the earlier cosine

e_h	Holding-share gap $\lvert s_e - s_c \rvert$	Amplification $\lvert \rho e_h - e_o \rvert^{-1}$	Mean $\hat{\tau}$	Sd $\hat{\tau}$
0.2	0.009	0.690	25.85	0.88
0.5	0.033	0.272	25.33	0.34
1.0	0.070	0.191	25.23	0.24
1.5	0.104	0.167	25.20	0.21
2.0	0.136	0.156	25.19	0.20
3.0	0.194	0.145	25.18	0.18

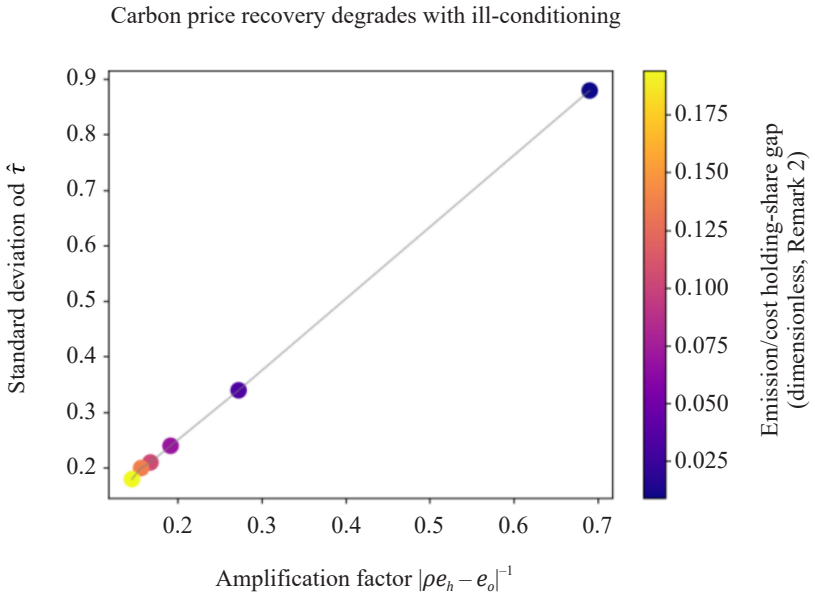


Figure 13. Dispersion of the recovered carbon price against the amplification factor, coloured by the dimensionless holding-share gap of Section 6.4 in place of the earlier informal cosine

dispersion of τ shrinks. When the two shares are nearly equal ($e_h = 0.20$, share gap 0.009), the amplification factor is largest and τ is most dispersed; as the gap widens the estimate tightens. This is the same qualitative message, recovery is best conditioned when emissions are structured differently from money, now stated with a statistic whose units and normalisation are explicit.

7.14 Bootstrap confidence interval coverage

Section 6.5 recommended a nonparametric bootstrap over records as the simplest way to obtain a variance estimate reflecting both stages of the estimator. We test the coverage of the resulting interval directly: for 300 independently generated panels ($N = 200$, baseline design, three per cent execution noise), we compute a nominal 95 per cent percentile bootstrap interval for $\hat{\tau}$ from 300 record-level bootstrap resamples of each panel, and record how often the interval contains the true $\tau = 25$. Table 24 reports the empirical coverage.

The plain percentile bootstrap under-covers by roughly fifteen percentage points, which is consistent with Section 6.5's diagnosis: the interval is centred on a biased $\hat{\rho}$ and does not correct for the skew that a near-zero denominator introduces into the sampling distribution of $\hat{\tau}$. This is the direct evidence for the paper's recommendation, throughout Sections 5.3, 6.5 and here, that a plain normal or percentile interval should not be trusted at face value for this estimator, that a bias-corrected-and-accelerated bootstrap or a Fieller-type set is preferable (Section 6.6), and that the full joint asymptotic treatment of Section 5.4 is the more durable long-run fix.

7.15 Comparison against a direct nonlinear least squares benchmark

No benchmark method has so far been compared against the two-stage estimator. We add the most natural alternative: a single-stage nonlinear least squares (NLS) fit that estimates (θ, ρ) jointly by minimising the sum of squared log-discrepancies between the observed cycle lengths and the cycle length the forward model of Section 4 predicts at each candidate (θ, ρ), without using the inventory-balance moment condition (9) at all. This benchmark uses only the optimality relationship, exactly the information a naïve "fit the cost curve" approach would use, and lets us ask whether the two-stage estimator's explicit use of the balance equation for θ is actually earning its keep. Table 25 compares the two estimators at the baseline design.

On the mechanically exact balance design of Section 7.2, the two-stage estimator dominates the NLS benchmark on both parameters, and by a wide margin on θ, exactly because it uses the balance equation directly rather than inferring deterioration only from the curvature of the optimality condition, which Section 5.1 already argued is a weak channel on its own. The two-stage estimator is also, in this implementation, roughly four orders of magnitude faster per panel (about 0.4 milliseconds against roughly twelve seconds for the NLS benchmark's iterative search), because it reduces to a one-dimensional search followed by a closed-form regression rather than a joint nonlinear optimisation over two parameters. This comparison is necessarily favourable to the two-stage estimator on the design where its extra moment condition (the balance equation) is exact; Section 7.5 shows that advantage narrows, and can be lost, once that moment condition is itself contaminated, so the fair reading of this benchmark is that the two-stage estimator's advantage is conditional on the balance equation carrying real information, not unconditional.

The comparison above uses the mechanically exact balance design throughout, which by construction favours the two-stage estimator, since it treats the balance equation

Table 24. Empirical coverage of a nominal 95% percentile bootstrap confidence interval for $\hat{\tau}$

Quantity	Value
Nominal coverage	95%
Empirical coverage (300 panels, 300 bootstrap resamples each)	80.3%
Mean interval width	2.23 (carbon price units)

Table 25. Two-stage estimator versus a direct nonlinear least squares benchmark, baseline design, three per cent execution noise

N	Two-stage RMSE($\hat{\theta}$)	NLS RMSE($\hat{\theta}$)	Two-stage RMSE($\hat{\rho}$)	NLS RMSE($\hat{\rho}$)
100	0.000	0.1749	0.1429	0.2579
200	0.000	0.0976	0.1285	0.1420

as noise-free information the joint benchmark cannot exploit any better than the two-stage method already does. A fairer comparison, with genuinely the same information set, also evaluates both estimators under the contaminated designs of Section 7.5, where the balance equation is no longer trustworthy and a method that does not treat it as exact might do better. We implemented a joint minimum-distance (MD) estimator that minimises a weighted combination of the (log) balance-equation residual and the optimality-condition residual simultaneously over (θ, ρ), rather than solving for θ first and treating it as fixed in the second stage. Table 26 compares its RMSE against the two-stage estimator's, on the mechanically exact design and on the two contamination levels of Table 15.

On the mechanically exact design the two-stage estimator remains at least as good as the joint benchmark on both parameters, consistent with Table 25's finding above: when the balance equation genuinely is exact, there is nothing for a joint estimator to gain by not trusting it fully, and the two-stage estimator's zero-variance first stage is, if anything, an advantage the joint method cannot replicate. The ranking reverses once the balance equation is contaminated: the joint minimum-distance estimator has materially lower RMSE for both θ and ρ at both contamination levels, because it does not force $\hat{\theta}$ to explain the balance equation exactly and so is not as fully exposed to the order-quantity and ending-inventory errors that dominate the two-stage estimator's first stage. The honest, and more useful, conclusion is accordingly conditional rather than a blanket ranking: the two-stage estimator is preferable when the analyst has verified that the balance equation is reliable (Section 7.5's diagnostic), and the joint estimator, or the external-θ workflow of Table 15, is preferable when that reliability cannot be established.

7.16 Misspecification and the value of the quantity channel

Two experiments probe what happens when the analyst uses the wrong model. Table 27 fits a classical order quantity that ignores deterioration to data in which items genuinely decay; the recovered cost ratio is biased downward, increasingly so as the deterioration rate rises. Table 28 contrasts the timing channel, which infers the rate from the shape of the optimality condition and collapses under noise, against the quantity channel, which reads the rate from the inventory balance and on the mechanically exact design recovers it essentially exactly. Section 7.5 is the necessary complement to Table 28: the quantity channel's advantage there is real but conditional on the balance being trustworthy, which real operational data will not always guarantee.

Table 26. Two-stage estimator against a joint minimum-distance/GMM benchmark with the same information set, mechanically exact and contaminated designs, N = 200, 80 replications per cell

Design	Two-stage RMSE($\hat{\theta}$)	Joint-MD RMSE($\hat{\theta}$)	Two-stage RMSE($\hat{\rho}$)	Joint-MD RMSE($\hat{\rho}$)
Mechanically exact balance	0.000	0.020	0.097	0.116
2% order-qty noise + 3% ending inv.	0.432	0.306	0.430	0.285
5% order-qty noise + 5% ending inv.	0.748	0.533	0.821	0.562

Table 27. Bias from ignoring deterioration and fitting a classical economic order quantity to perishable data

θ	True ρ	Correct $\hat{\rho}$	Naive $\hat{\rho}$	Naive bias %
0.05	18.868	18.779	18.695	–0.9
0.15	16.949	16.869	16.656	–1.7
0.25	15.385	15.311	15.008	–2.5
0.35	14.085	14.017	13.647	–3.1
0.45	12.987	12.925	12.506	–3.7
0.55	12.048	11.991	11.535	–4.3

Table 28. Mean absolute error (MAE) in the deterioration estimate, timing-only versus two-stage (mechanically exact balance design)

Demand spread	Timing only MAE	Two stage MAE
1.1	0.3	0
1.5	0.3	0
2.5	0.3	0
4.0	0.3	0
8.0	0.3	0

Table 29. One-at-a-time sensitivity of the recovered cost ratio and carbon price to twenty per cent perturbations in each generating parameter. Baseline ρ is 14.706 and τ is twenty-five

Parameter	$\hat{\rho}$ at –20%	$\hat{\rho}$ at +20%	$\hat{\tau}$ at –20%	$\hat{\tau}$ at +20%
K	12.291	16.973	25.64	25.48
h	15.546	13.819	25.48	25.60
c_d	15.449	13.896	25.49	25.59
θ	15.450	13.896	25.49	25.59
τ	15.406	13.983	20.47	30.61
e_o	14.046	15.217	25.47	25.62
e_h	16.047	13.445	25.69	25.45

8. Sensitivity analysis

The numerical study establishes that recovery works on clean and moderately noisy data and characterises its failure modes one at a time. Read together with Section 7.3's revised consistency result and Section 7.5's realistic-noise experiment, the sensitivity map below should be understood as bounding performance from above: it describes the best case in which the balance equation is exact, and Section 7.5 quantifies how much of that performance is lost once it is not.

Table 29 and Figure 14 perturb each generating parameter by twenty per cent in turn. The ordering cost dominates, moving the ratio across a range of nearly five units, followed by the holding emission intensity, the holding and spoilage costs and the carbon price itself. The two emission intensities have opposite effects, one raising the effective setup and the other raising the effective holding, which is the very tension that Theorem 2 shows is necessary for the carbon price to be identified at all.

The first error surface, in Table 30 and Figure 15, sweeps the cost ratio error over the panel size and the execution noise together. A noisy record can be partly offset by volume, though, per Section 7.3, only down to the bias floor at that noise level, not to zero.

The second error surface, in Table 31 and Figure 16, sweeps the deterioration error over the true rate and the measurement noise (again under the mechanically exact balance design; see the substantially larger errors in Section 7.5's Table 14 once the balance is realistically contaminated).

Taken together the three instruments draw a map that is real but narrower than a first look would suggest. The cost ratio is recovered reliably wherever the panel is not tiny, execution is not wild, and the balance equation is close to exact; the deterioration rate is recovered reliably under the same balance-exactness condition wherever the item genuinely decays; and the carbon price inherits the reliability of the ratio subject to the conditioning that Theorem 2 identifies, plus the sampling-distribution caveats of Section 6.5 whenever a point estimate rather than an interval is reported.

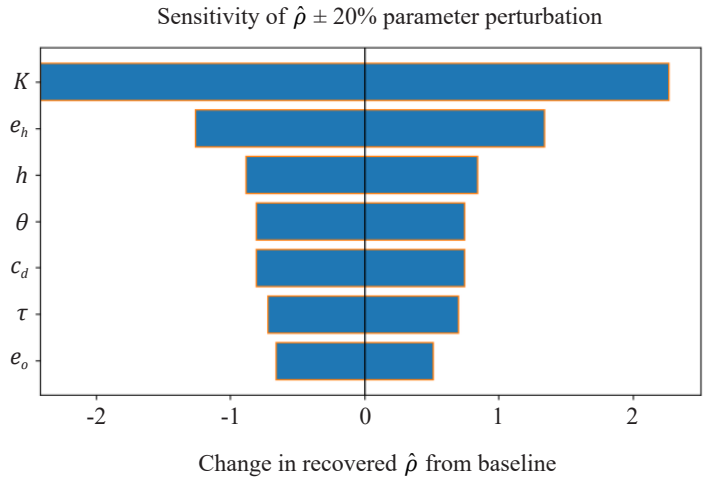


Figure 14. Change in the recovered cost ratio under a twenty per cent perturbation of each generating parameter

Table 30. Root mean squared error of the recovered cost ratio over sample size and execution noise

Noise %	$N = 20$	$N = 40$	$N = 80$	$N = 160$	$N = 320$
2	0.167	0.123	0.093	0.068	0.056
4	0.344	0.274	0.225	0.182	0.161
6	0.548	0.470	0.415	0.360	0.331
10	1.073	1.024	0.977	0.909	0.863
15	1.967	2.000	1.990	1.920	1.851

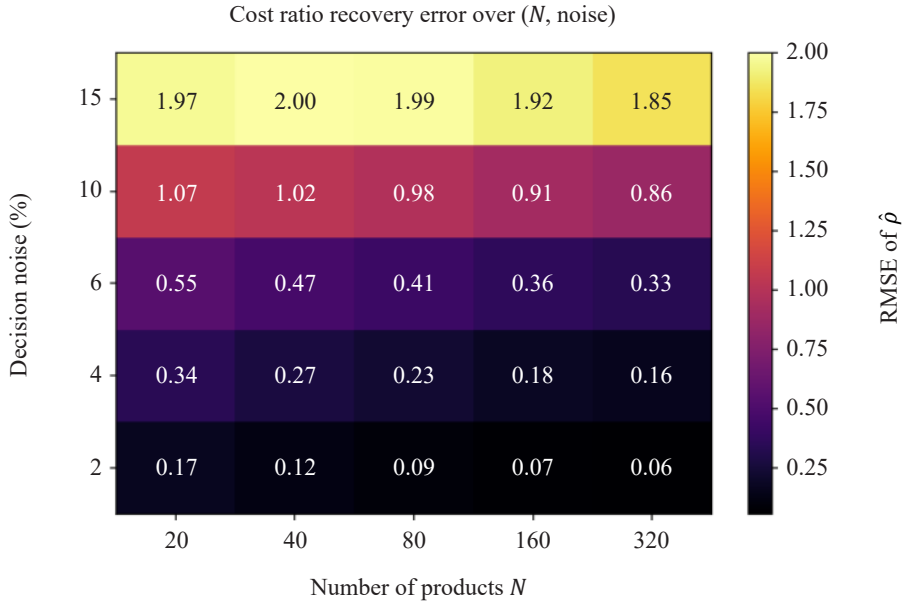


Figure 15. RMSE of the recovered cost ratio over the panel size and the execution noise (mechanically exact balance design)

Table 31. RMSE of the recovered deterioration rate over its true value and measurement noise (mechanically exact balance design)

Noise %	$\theta = 0.1$	$\theta = 0.2$	$\theta = 0.3$	$\theta = 0.4$	$\theta = 0.5$
1	0.0140	0.0148	0.0156	0.0164	0.0171
2	0.0309	0.0327	0.0345	0.0363	0.0380
4	0.0679	0.0810	0.0857	0.0902	0.0946
6	0.0870	0.1380	0.1568	0.1655	0.1738
10	0.0968	0.1893	0.2681	0.3303	0.3788

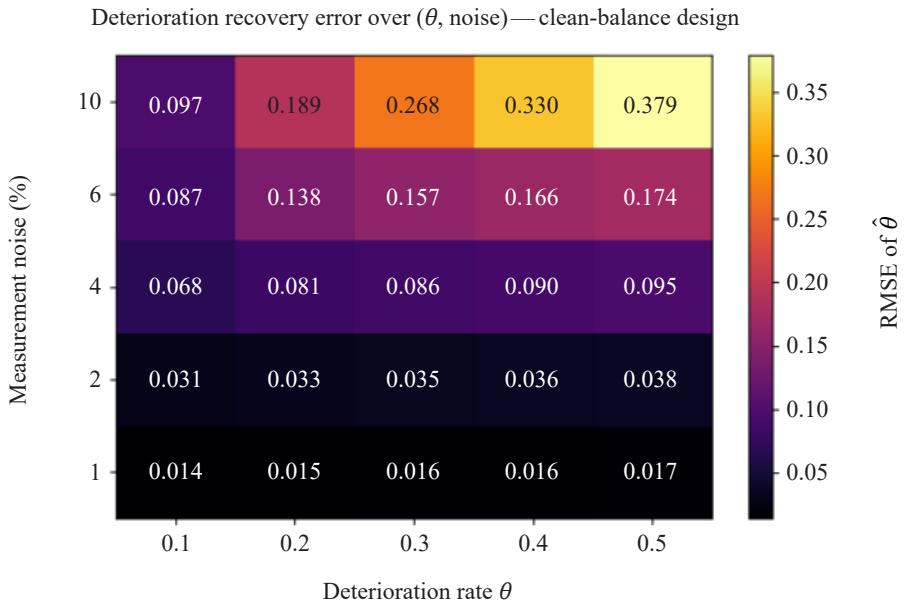


Figure 16. RMSE of the recovered deterioration rate over its true value and the measurement noise (mechanically exact balance design)

9. Managerial and policy implications

The framework turns a firm's own operational exhaust into a lens on its priorities, and several implications follow for the parties who might hold that lens. This section reflects the re-scoped reading of $\hat{\tau}$ established in Sections 5.6 and 6.4: a model-consistent shadow price, not a direct read-out of a firm's true carbon valuation.

For a regulator, $\hat{\tau}$ is a candidate input to, not a substitute for, an audit. A firm may report an internal carbon price, may publish a sustainability commitment, or may say nothing at all, and replenishment is a decision that the shadow price reads for consistency with the model of Section 4. A firm whose $\hat{\tau}$ is near zero is behaving, within the scope of that model, as if carbon were unpriced; a firm whose $\hat{\tau}$ matches the prevailing levy is behaving consistently with genuinely facing that levy. But Section 6.4 is explicit that $\hat{\tau}$ absorbs any omitted operational constraint that correlates with demand the way carbon would, so a low $\hat{\tau}$ is evidence consistent with weak internalisation, not proof of it, and a high $\hat{\tau}$ could in principle reflect an unmodelled cost rather than a carbon charge. We recommend that a regulator treat $\hat{\tau}$ as a flag that narrows where a conventional audit should look, cross-check it where possible against a firm's stated carbon price or against a period in which a carbon policy changed exogenously (which would shift $\hat{\tau}$ if genuine and leave it largely unchanged if the estimate were driven by an unrelated omitted cost), and rely on Theorem 2's conditioning number to know when the flag itself is statistically trustworthy; we deliberately avoid describing this as "a behavioural audit of genuine emission internalisation," reflecting this narrower, better-supported claim.

For an analyst benchmarking a competitor, the effective parameters are the durable objects. Individual costs cannot be separated from behaviour, as the confounding valley shows, so an analyst who claims to have reverse

engineered a rival's holding cost from its order pattern is overstating what the data support. What can be recovered, subject to the identification-unit caveat of Section 3.1, which matters as much for a rival's product mix as for one's own, is the effective cost-to-holding ratio and the deterioration rate, not the effective setup and the effective holding separately: by Proposition 2 the level of the pair (A, Ω) is not identified from decisions, only their ratio is, so it is ρ and θ, not A and Ω individually, that are often what matters for predicting how the rival will respond to a shock.

For a firm auditing itself, the residual diagnostic of Section 5.9 is a management signal. When the optimality fit is poor, relative to the threshold Section 5.9 now provides, the firm is either leaving money on the table or is not well described by a single cost structure, Section 7.8's pooling result shows the second possibility is not a corner case. The systematic direction of the bias under bounded rationality, documented in Section 7.7, is a caution that habitual satisficing does not merely add noise but distorts the picture in a predictable direction.

For the design of emission policy, the identification condition of Theorem 2 is a warning about instrument choice. A carbon price bites on behaviour only through the emission mix, and if that mix is proportional to the conventional cost mix the price is behaviourally invisible in this framework's terms, indistinguishable from a uniform cost increase. A policy that wishes its carbon signal to be legible in observable firm behaviour does well to attach the charge to the emission channel that is most distinct from ordinary costs, which for perishable trade is typically the holding side rather than the ordering side. We note that cold chain refrigeration is a holding linked, energy intensive emission source of the kind that Hu et al. [66] quantify as a major and reducible share of cold chain greenhouse gas emissions, and for which this condition may therefore be expected to hold more often than for ordering linked emissions, which is a reason for optimism about the framework's applicability to refrigerated perishable trade specifically, though not a reason to expect the same conditioning in categories where ordering-linked transport emissions dominate. Recent evidence that firms' internal carbon prices respond to external carbon constraints and need not equal regulated prices [57], and that internal pricing alone appears to be less effective than a carbon tax [58], reinforces the caution that a recovered shadow price should not be read as a declared internal price.

9.1 From recovered parameters to forward decisions and monitoring

The recovered quantities feed back into operations, which is where the value of the framework would be realised in deployment. By Equation (8), the optimal replenishment cycle for any demand the firm expects to face next is a function of the effective cost-to-holding ratio ρ and the deterioration rate θ alone; the effective setup and the effective holding are not separately required, and by Proposition 2 they are not separately available from decisions in any case. Because ρ and θ are the complete set of primitives the forward model requires, a planner who has recovered this pair can compute the optimal replenishment cycle directly, without needing, and without being able to obtain, A and Ω individually. We say "would be" deliberately: as Section 5.8 states, this paper specifies the computation and the workflow but has not implemented or field-tested the planning and monitoring loop it describes, and the accuracy inherited by any forward plan is bounded by the revised Section 7's findings, not by the more optimistic clean-design numbers alone.

Monitoring the shadow carbon price on a rolling window turns a single audit into a running indicator, and a control chart on that indicator can surface the moment a firm's replenishment stops looking consistent with the carbon price it professes, again read as a flag calling for corroboration, per the re-scoped interpretation above, rather than as a self-sufficient conclusion. The optimality-fit residual, now with the practical threshold of Section 5.9, doubles as an anomaly detector that points management to the lines where planning has drifted, or where the identification-unit assumption of Section 3.1 no longer holds.

10. Conclusion

We have taken a classical deterministic inventory model for a deteriorating item, embedded a carbon charge in it, and asked the inverse question that the forward literature never poses, namely what the observed replenishment reveals about the costs behind it. The answer rests on a reduction in which ordering, holding, spoilage and emission costs collapse into an effective setup and an effective holding, so that behaviour identifies a single ratio and the deterioration rate and nothing more, and, as Theorem 1 shows explicitly, a single exact record already identifies both, with additional records serving statistical rather than identificational purposes. On that foundation we built a two-stage estimator and defined a shadow carbon price that measures, within the scope of the model and its accounting inputs, how strongly a firm behaves as if emissions were costly.

The estimator is not consistent at fixed noise; it converges to a noise-dependent pseudo-true value, and we provide both the theoretical statement (Proposition 4) and a practical bootstrap bias correction (Section 5.5) that closes most of the resulting gap. The deterioration-rate channel is not robust once the inventory balance is not mechanically exact, and Section 7.5 quantifies how much of Table 17's accuracy is an artefact of the data-generating design rather than a property of the estimator. Pooling heterogeneous products produces a precise but

biased estimate, and Section 3.1 and Section 7.8 give both the diagnosis and practical mitigations. The carbon price is a model-consistent shadow price whose economic interpretation as a behavioural audit requires assumptions beyond the mathematics of Theorem 2, and Section 6.4 and Section 9 say so plainly. Statistical inference is supplied by Sections 5.4, 6.5, and 7.11, including the finding that a naïve bootstrap interval under-covers by roughly fifteen percentage points.

None of this overturns the paper's central contribution, the effective-parameter collapse, the identifiability theory built on it, and the demonstration that a genuinely informative shadow carbon price can be read from operational data when the emission signature is distinct from the conventional cost signature. What changes is the confidence with which each claim is made, and the tools an analyst is given to know when the framework's numbers should be trusted and when they should not.

10.1 Implementation in enterprise resource planning and inventory management systems

The framework's computational core is light enough to sit inside the systems firms already run: a one-dimensional search for the deterioration rate followed by a closed-form regression for the cost ratio, both of which run in well under a second over a routine catalogue, as Section 7.12's runtime comparison confirms. We are careful, however, about what has and has not been demonstrated. The paper specifies a computation and the data fields it needs, order sizes, sales rates, realised cycle lengths, and accounting costs, all standard fields in an enterprise resource planning or inventory management database, and describes how a scheduled batch job could expose the recovered rate, ratio and shadow carbon price to a planning module, a sustainability dashboard, and a drift-detection flag. It does not implement, benchmark on real transaction data, or field-test such a deployment, and the phrase "decision support system" overstates this: we use "engine", "module" and "workflow" throughout to describe what has actually been built and validated, and reserve "system" for the deployed artefact that a future implementation study would need to construct and evaluate against real operational data, ideally including a comparison to the current planning practice it would replace.

10.2 Limitations

The framework rests on assumptions that bound its reach, and this section restates the limitations already present and adds further ones. It reads behaviour through a deterministic deteriorating item model with constant demand within a cycle and no shortages, so a firm that carries safety stock against uncertain demand or that rations under scarcity departs from the premise the recovery inverts. The optimality assumption is load bearing, and the bounded rationality experiment shows that habitual satisficing biases the recovered costs in a fixed direction rather than merely adding noise. The carbon recovery depends on accounting costs supplied from outside the model, is well conditioned only when the emission signature differs from the conventional cost signature, and, as Section 6.4 stresses, measures a model-consistent shadow price rather than a directly validated internal carbon valuation.

The two-stage estimator is not consistent at fixed noise, only in a joint asymptotic regime; its joint asymptotic distribution, including the second-stage regression's generated-regressor variance, is now derived in closed form in Section 5.4 and validated numerically, though the first-order approximation there does not capture the second-order bias term of Appendix D, which must be tracked separately (Sections 5.5, 5.4, 6.5). The deterioration-rate channel recovers θ essentially exactly only because the data-generating design preserves the inventory balance mechanically; Section 7.5 shows materially larger errors once order-quantity measurement error and nonzero ending inventory are introduced, and the clean-design numbers should not be read as representative of achievable field accuracy. The common-parameter assumption (Assumption 4) is not automatically satisfied by a panel of records and must be checked, ideally through the identification-unit reasoning of Section 3.1, before pooling; Section 7.8 shows what pooling across a genuine mixture costs. Confidence intervals built from a plain bootstrap under-cover materially, and we recommend bias-corrected-and-accelerated bootstrap or Fieller-type intervals in their place (Section 6.6). Finally, the demand rate is treated as known context rather than estimated jointly, which keeps the identification clean at the cost of assuming the sales history is measured well, and Section 7.4 finds that this is precisely the input whose mismeasurement the carbon-price estimate is least able to absorb. We could not validate the framework against real or semi-real transaction data in this study, for lack of access to a suitable licensed or public dataset with matched shelf-life and accounting/emission records; this is the single highest-priority gap for future work to close, ahead of the extensions listed below.

10.3 Future work

Several directions follow directly from the results above rather than being generic suggestions. The most pressing item this revision closes rather than leaving open: a full joint asymptotic treatment of the two-stage estimator, casting both stages as stacked estimating equations, deriving the GMM sandwich variance with the first-stage Jacobian propagated through by the delta method, and characterising the small-sample behaviour of Fieller-type confidence sets for τ near the Theorem 2 boundary, is now given in Section 5.4 and Section 6.6 respectively, validated numerically against Monte Carlo simulation in Table 3 and Table 7. What remains open

beyond that derivation is as follows. A second direction is a hierarchical or random-coefficients extension that estimates a distribution of (θ, ρ) across products explicitly, rather than requiring the analyst to pre-cluster into homogeneous identification units by hand as Section 3.1 currently recommends. A third is a formal treatment of measurement error on the order quantity and on ending inventory within the estimator itself, an errors-in-variables correction for Equation (9) rather than only characterising its damage empirically as Section 7.5 does; a control-variate or instrumental approach in the spirit of the errors-in-variables literature is a natural next step. A fourth is to extend the model to temperature-dependent or age-dependent deterioration building on the classification of deterioration forms in Bakker et al. [17] and on cold chain energy and emission studies such as Hu et al. [66], which would let the framework speak more directly to cold-chain applications. A fifth is estimation that is robust to outliers and to a mixture of satisficing and optimising firms in the same panel, for which the noisy data treatment of Aswani et al. [4] and the contaminated design comparison of Section 7.12 are a starting point rather than a solution. And the most consequential for the paper's stated motivation is empirical application to real replenishment records together with independent evidence on the firms' actual carbon prices or policies, which is the only way to move from "model-consistent shadow price" back toward the stronger "behavioural audit" language, which this paper deliberately sets aside.

Appendix A. Derivation of the forward relations

The cycle dynamics (1) are linear with constant coefficients. Writing the equation as $\dot{I} + \theta I = -D$ and applying the integrating factor $e^{\theta t}$ gives $\frac{d}{dt}[e^{\theta t} I] = -De^{\theta t}$. Integrating from t to T and imposing $I(T) = 0$ yields $I(t) = \left(\frac{D}{\theta}\right)[e^{\theta(T-t)} - 1]$, which is (2), and evaluating at $t = 0$ gives the order quantity. The holding integral follows from $\int_0^T I(t)dt = \left(\frac{D}{\theta}\right)\int_0^T [e^{\theta(T-t)} - 1]\,dt = \left(\frac{D}{\theta}\right)\left[\frac{(e^{\theta T} - 1)}{\theta} - T\right]$ $= \left(\frac{D}{\theta^2}\right)[e^{\theta T} - 1 - \theta T]$. The spoilage is the order less the quantity sold, $Q - DT$, and substituting the order quantity and factoring reproduces the final identity of (3), in which spoilage equals the deterioration rate times the holding integral.

Appendix B. Effective collapse and optimality

Dividing the cycle cost by T gives the three term rate in (4). The holding integral per unit time is $D\widehat{H}_b(T; \theta)$ with $\widehat{H}_b$ as defined, and the spoilage per unit time is $\theta D\widehat{H}_b(T; \theta)$ by the identity of Appendix A. Grouping the coefficient of $\widehat{H}_b$ collects $h + \tau e_h$ from holding and emission with $c_d\theta$ from spoilage, which yields the effective holding $\Omega = h + c_d\theta + \tau e_h$, while the two terms carried by $1/T$ collect into the effective setup $A = K + \tau e_o$. This is (6).

Differentiating, $\widehat{H}_b^{\prime(T;\theta)} = \frac{[\theta(e^{\theta T} - 1)T - (e^{\theta T} - 1 - \theta T)]}{(\theta^2 T^2)}$, which is positive for $T > 0$ because the numerator vanishes at $T = 0$ with a positive derivative and grows thereafter. The cost rate has second derivative $\frac{2A}{T^3} + \Omega D\widehat{H}_b^{\prime\prime(T;\theta)}$, and since $\widehat{H}_b$ is convex in T the whole expression is positive, so the stationary point of (7) is the unique global minimum. The small rate limit uses $e^{\theta T} = 1 + \theta T + \theta^2 T^2/2 + O(\theta^3)$, whence $\widehat{H}_b \to \frac{T}{2}$ and the optimality condition reduces to the square root formula $T^* = \sqrt{\frac{2A}{(\Omega D)}}$.

This appendix also underlies Lemma 1 in Section 6: the same series expansion applied to the balance function $g(\theta) = \left(\frac{D}{\theta}\right)(e^{\theta T} - 1)$ shows every term of $\frac{g(\theta)}{D} = T + \frac{\theta T^2}{2} + \frac{\theta^2 T^3}{6} + \cdots$ is non-negative and strictly positive for $\theta > 0$, so g is strictly increasing on $[0, \infty)$ for fixed D, $T > 0$, which is the monotonicity property that gives single-record identification of θ.

Appendix C. The inverse estimators

For the deterioration rate, the balance residual in (9) is a smooth scalar objective whose minimiser is found by a one-dimensional search, and its curvature at the optimum is bounded away from zero whenever the spoilage surplus responds to the rate, which holds as soon as the deterioration exposure is not negligible. For the cost ratio, the optimality condition (8) states $D_i\widehat{H}_b^{\prime(T_i;\theta)} = \frac{\rho}{T_i^2}$, so with the constructed quantities of (10) the record equations are $x_{1i} = \rho x_{2i}$. The ordinary least squares solution minimising $\sum_i (x_{1i} - \rho x_{2i})^2$ is $\hat{\rho} = \frac{\sum_i x_{1i} x_{2i}}{\sum_i x_{2i}^2}$, which is (11). For the carbon price, substituting $A = K + \tau e_o$ and $\Omega = h + c_d\theta + \tau e_h$ into $\rho = A/\Omega$ and solving for τ gives (20) whenever $\rho e_h \neq e_o$.

As Section 5.3 discusses, (11) should not be read as a textbook OLS estimator: x_{1i} depends on the first-stage estimate $\hat{\theta}$ rather than the true θ, and both x_{1i} and x_{2i} are

nonlinear functions of the same noisy T_i, so the usual OLS variance formula understates the true sampling variance and does not capture the generated-regressor and errors-in-variables components discussed there, now formally corrected by the sandwich variance of Section 5.4. Appendix D's second-order expansion is the source of the resulting bias.

Appendix D. Small noise bias of the ratio estimator

Let the observed cycle be $T_i = T_i^* e^{\sigma\varepsilon_i}$ with ε_i symmetric of unit variance and σ the noise level. The denominator regressor is $x_{2i} = T_i^{-2} = T_i^{*-2} e^{-2\sigma\varepsilon_i}$, with expectation $T_i^{*-2}[1 + 2\sigma^2 + O(\sigma^4)]$ under symmetry. The numerator regressor is $x_{1i} = D_i \widehat{H}_b^{\prime(T_i;\theta)}$, and a second order Taylor expansion of $\widehat{H}_b'$ about T_i^* gives

$$\mathrm{E}[x_{1i}] = D_i \widehat{H}_b^{\prime(T_i^*;\theta)} + 1/2 D_i T_i^{*2} \widehat{H}_b^{\prime\prime\prime(T_i^*;\theta)} + O(\sigma^4),$$

where the first order term vanishes in expectation by symmetry. Since $D_i \widehat{H}_b'(T_i^*; \theta) = \rho T_i^{*-2}$ at the truth, the ratio of expectations departs from ρ by a term proportional to σ^2 whose sign is set by the third derivative of $\widehat{H}_b'$ weighted across records. This is exactly the pseudo-true bias $b(\sigma)$ of the revised Proposition 4 in Section 7.3: it is of second order in the noise and vanishes as $\sigma \to 0$, in agreement with the exact recovery observed on clean data, but at any fixed $\sigma > 0$ it survives as $N \to \infty$ because it is a bias in the expectation of the constructed regressors, not a sampling fluctuation that averages away. The sampling variance of $\hat{\rho}$ around this biased centre is the variance of a ratio of averages of the constructed quantities and scales as N^{-1}, giving the $N^{-1/2}$ law for the standard deviation that Table 11 confirms empirically for the dispersion component, while the bias floor confirms the systematic component derived here.

Declaration of generative AI use

The author declares that no generative AI tools were used in the writing, data analysis, or generation of figures for this manuscript.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial or not for profit sectors.

Data availability

All data in this study are synthetic and were generated by Monte Carlo simulation; no empirical, proprietary or human subject data were used. The data generating designs, parameter values, sample sizes, noise levels and numbers of replications are stated in full in Sections 7 and 8, so that every table and figure can be regenerated from those specifications. The simulation code and the generated datasets are not yet deposited in a public repository because the code base is being extended in ongoing follow-up work by the author. They are available from the corresponding author (janardanbeheragreetsyou@gmail.com) upon reasonable request, without restriction on non-commercial academic use, and will be deposited in a public GitHub repository with a permanent Zenodo DOI upon completion of that work.

Conflict of interest

The author declares that there are no known financial or personal relationships or interests that could have appeared to influence the work reported in this paper.

References

[1] Ahuja RK, Orlin JB. Inverse optimization. *Operations Research*. 2001; 49(5): 771-783. doi:10.1287/opre.49.5.771.10607.

[2] Iyengar G, Kang W. Inverse conic programming with applications. *Operations Research Letters*. 2005; 33(3): 319-330. doi:10.1016/j.orl.2004.04.007.

[3] Keshavarz A, Wang Y, Boyd S. Imputing a convex objective function. In: *2011 IEEE International Symposium on Intelligent Control*. Denver, CO, USA: IEEE; 2011. p.613-619. doi:10.1109/ISIC.2011.6045410.

[4] Aswani A, Shen ZJM, Siddiq A. Inverse optimization with noisy data. *Operations Research*. 2018; 66(3): 870-892. doi:10.1287/opre.2017.1705.

[5] Mohajerin Esfahani P, Shafieezadeh-Abadeh S, Hanasusanto GA, Kuhn D. Data-driven inverse optimization with imperfect information. *Mathematical Programming*. 2018; 167(1): 191-234. doi:10.1007/s10107-017-1216-6.

[6] Dong C, Chen Y, Zeng B. Generalized inverse optimization through online learning. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. Red Hook, NY, United States: Curran Associates Inc.; 2018. p.86-95.

[7] Chan TCY, Kaw N. Inverse optimization for the recovery of constraint parameters. *European Journal of Operational Research*. 2020; 282(2): 415-427.

doi:10.1016/j.ejor.2019.09.027.
[8] Saez-Gallego J, Morales JM, Zugno M, Madsen H. A data-driven bidding model for a cluster of price-responsive consumers of electricity. *IEEE Transactions on Power Systems*. 2016; 31(6): 5001-5011. doi:10.1109/TPWRS.2016.2530843.
[9] Chan TCY, Mahmood R, Zhu IY. Inverse optimization: theory and applications. *Operations Research*. 2023; 73(2): 1046-1074. doi:10.1287/opre.2022.0382.
[10] Harris FW. How many parts to make at once. *Factory, the Magazine of Management*. 1913; 10(2): 135-136.
[11] Wagner HM, Whitin TM. Dynamic version of the economic lot size model. *Management Science*. 1958; 5(1): 89-96. doi:10.1287/mnsc.5.1.89.
[12] Cárdenas-Barrón LE, Chung KJ, Treviño-Garza G. Celebrating a century of the economic order quantity model in honor of Ford Whitman Harris. *International Journal of Production Economics*. 2014; 155: 1-7. doi:10.1016/j.ijpe.2014.07.002.
[13] Ghare PM, Schrader GF. A model for exponentially decaying inventory. *Journal of Industrial Engineering*. 1963; 14(5): 238-243.
[14] Covert RP, Philip GC. An EOQ model for items with Weibull distribution deterioration. *AIIE Transactions*. 1973; 5(4): 323-326.
[15] Nahmias S. Perishable inventory theory: a review. *Operations Research*. 1982; 30(4): 680-708. doi:10.1287/opre.30.4.680.
[16] Goyal SK, Giri BC. Recent trends in modeling of deteriorating inventory. *European Journal of Operational Research*. 2001; 134(1): 1-16. doi:10.1016/S0377-2217(00)00248-4.
[17] Bakker M, Riezebos J, Teunter RH. Review of inventory systems with deterioration since 2001. *European Journal of Operational Research*. 2012; 221(2): 275-284. doi:10.1016/j.ejor.2012.03.004.
[18] Sarkar B. An EOQ model with delay in payments and time varying deterioration rate. *Mathematical and Computer Modelling*. 2012; 55(3-4): 367-377. doi:10.1016/j.mcm.2011.08.009.
[19] Hua G, Cheng TCE, Wang S. Managing carbon footprints in inventory management. *International Journal of Production Economics*. 2011; 132(2): 178-185. doi:10.1016/j.ijpe.2011.03.024.
[20] Benjaafar S, Li Y, Daskin M. Carbon footprint and the management of supply chains: insights from simple models. *IEEE Transactions on Automation Science and Engineering*. 2013; 10(1): 99-116. doi:10.1109/TASE.2012.2203304.
[21] Chen X, Benjaafar S, Elomri A. The carbon-constrained EOQ. *Operations Research Letters*. 2013; 41(2): 172-179. doi:10.1016/j.orl.2012.12.003.
[22] Toptal A, Özlü H, Konur D. Joint decisions on inventory replenishment and emission reduction investment under different emission regulations. *International Journal of Production Research*. 2014; 52(1): 243-269. doi:10.1080/00207543.2013.836615.
[23] Battini D, Persona A, Sgarbossa F. A sustainable EOQ model: theoretical formulation and applications. *International Journal of Production Economics*. 2014; 149: 145-153. doi:10.1016/j.ijpe.2013.06.026.
[24] He P, Zhang W, Xu X, Bian Y. Production lot-sizing and carbon emissions under cap-and-trade and carbon tax regulations. *Journal of Cleaner Production*. 2015; 103: 241-248. doi:10.1016/j.jclepro.2014.08.102.
[25] Dye CY, Yang CT. Sustainable trade credit and replenishment decisions with credit-linked demand under carbon emission constraints. *European Journal of Operational Research*. 2015; 244(1): 187-200. doi:10.1016/j.ejor.2015.01.026.
[26] Afriat SN. The construction of utility functions from expenditure data. *International Economic Review*. 1967; 8(1): 67-77.
[27] Varian HR. The nonparametric approach to demand analysis. *Econometrica*. 1982; 50(4): 945-973.
[28] Lewbel A. The identification zoo: meanings of identification in econometrics. *Journal of Economic Literature*. 2019; 57(4): 835-903. doi:10.1257/jel.20181361.
[29] den Boer AV. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*. 2015; 20(1): 1-18. doi:10.1016/j.sorms.2015.03.001.
[30] Gallego G, van Ryzin G. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*. 1994; 40(8): 999-1020. doi:10.1287/mnsc.40.8.999.
[31] Power DJ, Sharda R. Model-driven decision support systems: concepts and research directions. *Decision Support Systems*. 2007; 43(3): 1044-1061. doi:10.1016/j.dss.2005.05.030.
[32] Arnott D, Pervan G. Eight key issues for the decision support systems discipline. *Decision Support Systems*. 2008; 44(3): 657-672. doi:10.1016/j.dss.2007.09.003.
[33] Holsapple C, Lee-Post A, Pakath R. A unified foundation for business analytics. *Decision Support Systems*. 2014; 64: 130-141. doi:10.1016/j.dss.2014.05.013.
[34] Bertsimas D, Kallus N. From predictive to prescriptive analytics. *Management Science*. 2019; 66(3): 1025-1044. doi:10.1287/mnsc.2018.3253.
[35] Elmachtoub AN, Grigas P. Smart "predict, then optimize". *Management Science*. 2021; 68(1): 9-26. doi:10.1287/mnsc.2020.3922.
[36] Bodur M, Chan TCY, Zhu IY. Inverse mixed integer optimization: polyhedral insights and trust region methods. *INFORMS Journal on Computing*. 2022; 34(3): 1471-1488. doi:10.1287/ijoc.2021.1138.
[37] Galli L, Levato T, Schoen F, Tigli L. Prescriptive analytics for inventory management in health care. *Journal of the Operational Research Society*. 2021; 72(10): 2211-2224. doi:10.1080/01605682.2020.177

6167.
[38] Hansen LP. Large sample properties of generalized method of moments estimators. *Econometrica*. 1982; 50(4): 1029-1054.
[39] Newey WK. A method of moments interpretation of sequential estimators. *Economics Letters*. 1984; 14(2-3): 201-206.
[40] Murphy KM, Topel RH. Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics*. 1985; 3(4): 370-379.
[41] Newey WK, McFadden D. Large sample estimation and hypothesis testing. In: Engle RF, McFadden DL. (eds.) *Handbook of Econometrics*. Vol. 4. Amsterdam: Elsevier; 1994. p.2111-2245. doi:10.1016/S1573-4412(05)80005-4.
[42] Fieller EC. Some problems in interval estimation. *Journal of the Royal Statistical Society, Series B*. 1954; 16(2): 175-185. doi:10.1111/j.2517-6161.1954.tb00159.x.
[43] Efron B. Better bootstrap confidence intervals. *Journal of the American Statistical Association*. 1987; 82(397): 171-185. doi:10.1080/01621459.1987.10478410.
[44] Efron B, Tibshirani RJ. *An Introduction to the Bootstrap*. New York: Chapman & Hall; 1993.
[45] Besbes O, Fonseca Y, Lobel I. Contextual inverse optimization: offline and online learning. *Operations Research*. 2025; 73(1): 424-443. doi:10.1287/opre.2021.0369.
[46] Zattoni Scroccaro P, Atasoy B, Mohajerin Esfahani P. Learning in inverse optimization: incenter cost, augmented suboptimality loss, and algorithms. *Operations Research*. 2024; 73(5): 2661-2679. doi:10.1287/opre.2023.0254.
[47] Gupta R, Zhang Q. Decomposition and adaptive sampling for data-driven inverse linear optimization. *INFORMS Journal on Computing*. 2022; 34(5): 2720-2735. doi:10.1287/ijoc.2022.1162.
[48] Crönert T, Martin L, Minner S, Tang CS. Inverse optimization of integer programming games for parameter estimation arising from competitive retail location selection. *European Journal of Operational Research*. 2024; 312(3): 938-953. doi:10.1016/j.ejor.2023.06.041.
[49] Huang G, Yu G, Zhao X. Refining data-driven upfront reservation discount pricing via inverse inferring newsvendor transactions. *Production and Operations Management*. 2026; 35(5): 1766-1784. doi:10.1177/10591478251391628.
[50] Erkip NK. Can accessing much data reshape the theory? Inventory theory under the challenge of data-driven systems. *European Journal of Operational Research*. 2023; 308(3): 949-959. doi:10.1016/j.ejor.2022.08.024.
[51] Paul A, Pervin M, Roy SK, Maculan N, Weber GW. A green inventory model with the effect of carbon taxation. *Annals of Operations Research*. 2022; 309(1): 233-248. doi:10.1007/s10479-021-04143-8.
[52] Jain R, Mittal M, Mangla SK, Baraiya R. Optimizing supply chain strategies for deteriorating items and imperfect manufacturing under carbon emission regulations. *Computers & Industrial Engineering*. 2023; 182: 109350. doi:10.1016/j.cie.2023.109350.
[53] Mahato F, Choudhury M, Das S, Mahata GC. Optimal pricing and replenishment decisions for non-instantaneous deteriorating items with a fixed lifetime and partial backordering under carbon regulations. *Environment, Development and Sustainability*. 2024; 26(8): 21399-21432. doi:10.1007/s10668-023-03536-y.
[54] Pervin M. A sustainable deteriorating inventory model with backorder and controllable carbon emission by using green technology. *Environment, Development and Sustainability*. 2025; 27(10): 25005-25041. doi:10.1007/s10668-024-04717-z.
[55] Bhadoriya A, Jani MY, Shukla M, Abbas M. Is social, economic, and environmental sustainability essential for inventory modelling? *RAIRO-Operations Research*. 2025; 59(4): 2177-2200. doi:10.1051/ro/2025066.
[56] Konstantaras I, Skouri K, Benkherouf L. Optimizing inventory decisions for a closed-loop supply chain model under a carbon tax regulatory mechanism. *International Journal of Production Economics*. 2021; 239: 108185. doi:10.1016/j.ijpe.2021.108185.
[57] Trinks A, Mulder M, Scholtens B. External carbon costs and internal carbon pricing. *Renewable and Sustainable Energy Reviews*. 2022; 168: 112780. doi:10.1016/j.rser.2022.112780.
[58] Zakriya M, Ansaram K, Mazza P. The race for carbon pricing among firms. *International Review of Financial Analysis*. 2026; 109: 104742. doi:10.1016/j.irfa.2025.104742.
[59] Zhang YJ, Sun YF, Huo BF. The optimal product pricing and carbon emissions reduction profit allocation of CET-covered enterprises in the cooperative supply chain. *Annals of Operations Research*. 2023; 329(1): 871-899. doi:10.1007/s10479-021-04162-5.
[60] Hansen BE, Lee S. Inference for iterated GMM under misspecification. *Econometrica*. 2021; 89(3): 1419-1447. doi:10.3982/ECTA16274.
[61] Dufour JM, Flachaire E, Khalaf L, Zalghout A. Identification-robust methods for comparing inequality with an application to regional disparities. *The Journal of Economic Inequality*. 2024; 22(2): 433-452. doi:10.1007/s10888-023-09600-x.
[62] Ratsimalahelo Z. Re-examining confidence intervals for ratios of parameters. *Econometrics*. 2025; 13(3): 37. doi:10.3390/econometrics13030037.
[63] Cavaliere G, Gonçalves S, Nielsen MØ, Zanelli E. Bootstrap inference in the presence of bias. *Journal of the American Statistical Association*. 2024; 119(548): 2908-2918. doi:10.1080/01621459.2023.2

284980.
[64] Mandi J, Kotary J, Berden S, Mulamba M, Bucarey V, Guns T, et al. Decision-focused learning: foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research*. 2024; 80: 1623-1701. doi:10.1613/jair.1.15320.
[65] Sadana U, Chenreddy A, Delage E, Forel A, Frejinger E, Vidal T. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*. 2025; 320(2): 271-289. doi:10.1016/j.ejor.2024.03.020.
[66] Hu G, Mu X, Xu M, Miller SA. Potentials of GHG emission reductions from cold chain systems: case studies of China and the United States. *Journal of Cleaner Production*. 2019; 239: 118053. doi:10.1016/j.jclepro.2019.118053.