The Bullwhip Effect in Closed-Loop Supply Chains with Periodic-Review (t, S) Policies and Mandatory Refurbishment Constraints: A Controlled Comparison with the Classical Linear Supply Chain

Rico Wojanowski[1]

[1]Department of Supply Chain Management, Hochschule Darmstadt, Darmstadt, Germany

Abstract

Mandatory refurbishment quotas are increasingly imposed on manufacturers as part of circular-economy policy, yet their effect on the bullwhip phenomenon under realistic periodic-review inventory control is not well understood, and the existing closed-loop literature reports conflicting signs for the effect of return flows on order variability. We develop a discrete-event Monte-Carlo model of a four-echelon supply chain (retailer, wholesaler, producer, refurbishment station) operating a periodic-review (t, S) policy in which the producer must source a mandated fraction ρ of its build from refurbished components; when refurbished stock is insufficient, production is bindingly constrained. The central design choice of this study is that the *same* model reproduces the classical linear (forward) supply chain exactly, by setting $\rho = 0$: this yields a fully controlled comparison in which the demand process, ordering policy, lead times and review intervals are held identical, and the *only* difference between the two systems is the mandatory-refurbishment coupling at the producer.

We measure the bullwhip effect using the mean-normalised variance ratio computed directly from simulated order streams, separate the periodic-review batching component from constraint-induced effects, and contrast two ordering policies (classical order-up-to, OUT; and proportional order-up-to, POUT, Dejonckheere et al., 2003) across a full-factorial design of 576 configurations replicated 30 times (34,560 runs). All results are reported with 95% bootstrap confidence intervals.

The controlled comparison yields a clear and, at first sight, counter-intuitive result. Relative to the linear chain, the closed-loop constraint leaves the downstream echelons essentially unchanged (amplification ratio $AR \approx 1.00$ at retailer and wholesaler) but *dampens* producer bullwhip ($AR < 1$), and the damping strengthens as the constraint binds more tightly—falling to $AR \approx 0.77$ (OUT) and 0.91 (POUT) at a low refurbishment yield. We trace the mechanism to the truncation of the upper tail of the producer order distribution: a binding refurbishment availability caps the large catch-up production pulses, so the constraint behaves as a low-pass filter on the producer's order signal. We further show that the ordering policy is the dominant determinant of absolute bullwhip levels (an order-of-magnitude effect, far larger than any single physical parameter), that the periodic-review interval is the single largest amplifier of producer variability, and that system feasibility is governed by refurbishment yield ϕ through a mass-balance condition $\phi \cdot \gamma_{ret} \geq 1$. Finally, a policy-dependent trade-off emerges: the variance-damping POUT rule operates at a substantially higher constraint frequency, because smoother ordering accumulates less refurbished buffer. The findings give an explicit mechanism for the recurring observation that closing a supply chain can reduce

variability, and reposition ordering-policy choice as the primary operational lever for stable circular supply chains.

Keywords: Bullwhip effect; Closed-loop supply chain; Linear supply chain; Periodic review; (t, S) policy; Proportional order-up-to; Refurbishment constraints; Circular economy; Extended producer responsibility

Contents

1 Introduction **4**

1.1 Motivation 4

1.2 The bullwhip effect and its closed-loop puzzle 4

1.3 Approach: a controlled in-model comparison 5

1.4 Contributions 5

1.5 Structure of the paper 6

2 Related Work and Positioning **6**

2.1 Foundations of the bullwhip effect 6

2.2 Periodic review and the control-theoretic view 7

2.3 The bullwhip effect in closed-loop supply chains 8

2.4 Circular-economy strategies and the policy context 9

3 Model **9**

3.1 Physical structure 10

3.2 Demand and forecasting 11

3.3 Inventory policy 11

3.4 Refurbishment constraint and the linear chain 12

3.5 Steady-state feasibility condition 13

3.6 Why variance is measured, not propagated 13

3.7 Bullwhip metric and normalisation 14

3.8 Sequence of events 14

4 Experimental Design **15**

4.1 A worked single-period illustration 16

5 Results **16**

5.1 The controlled comparison: localised damping 16

5.2 The mechanism: a low-pass filter on producer orders 18

5.3 Absolute levels: ordering policy and review interval 19

5.4 Interaction effects across the full design 20
5.5 Feasibility is governed by yield, not delay 21
5.6 A policy-dependent trade-off 22
5.7 Robustness and internal validation 22
5.8 Factor hierarchy across the full design 24
5.9 How the damping interacts with the review interval 26
5.10 Summary of results 27

6 Sensitivity Analysis **28**
6.1 Proportional controller horizon 29
6.2 Demand variability 29
6.3 Smoothing parameter and the closed-loop contrast 29
6.4 Refurbishment delay revisited 30
6.5 Summary of sensitivity 30

7 Discussion **31**
7.1 Does the constraint amplify the bullwhip effect? 31
7.2 The constraint as a filter, not a destabiliser 32
7.3 Ordering policy as the primary operational lever 32
7.4 Yield-governed feasibility and its managerial reading 33
7.5 Positioning against specific prior findings 33
7.6 Managerial and policy implications 34
7.7 Scope and limitations 34
7.8 Implications for research and practice 35

8 Conclusion **35**

Reproducibility **36**

A Simulation algorithm **40**

B Per-factor marginal results **40**

1 Introduction

1.1 Motivation

The transition toward a circular economy is reshaping the operational logic of manufacturing supply chains. Regulatory instruments increasingly oblige producers to recover, refurbish and reuse a defined share of their material inputs, and the regulatory perimeter is widening: in the European Union, extended producer responsibility (EPR) schemes already cover packaging, batteries, vehicles and electrical and electronic equipment, with textiles being added, while the Waste Electrical and Electronic Equipment (WEEE) Directive sets binding collection and recovery targets for producers (Batista et al., 2018). Despite two decades of such policy, recovery performance remains far below target—only about 40% of WEEE is recycled in the EU and only a fraction of facilities apply high-quality treatment—so the pressure on producers to internalise larger recovered-material quotas is increasing rather than easing.

These mandates introduce a supply-side coupling that is absent from classical forward supply chains: production is no longer limited only by demand and capacity, but also by the availability of successfully refurbished components. When that availability falls short of the mandated quota, output is bindingly constrained. The operational question that follows is the subject of this paper: does such a binding refurbishment coupling make a supply chain *more* or *less* volatile than the equivalent forward chain—and through what mechanism?

1.2 The bullwhip effect and its closed-loop puzzle

The bullwhip effect—the upstream amplification of demand variability—has been studied since Forrester (1961), with Lee et al. (1997) identifying four operational causes (demand-signal processing, rationing and gaming, order batching, and price fluctuation) and Chen et al. (2000) quantifying the amplifying role of exponential-smoothing forecasts under continuous-review order-up-to control. Real inventory systems, however, are typically reviewed periodically rather than continuously, and Dejonckheere et al. (2003) showed, using a control-theoretic formulation, that the review interval and the structure of the ordering rule are themselves decisive determinants of variance propagation—to the point that the classical order-up-to rule *guarantees* amplification regardless of the forecasting method used.

In closed-loop systems the picture is more complicated and, in the published record, contradictory. Returns can *dampen* ordering, because a recovered unit substitutes for a unit that would otherwise be procured new, reducing the order volume; or they can *tighten* production, because a binding recovery quota that cannot be met from available returns forces the producer into catch-up. Zhou and Disney (2006) find that returns may amplify or dampen the bullwhip effect depending on return rate and lead time; Dominguez et al. (2020) show that remanufacturing lead-time *variability* damages dynamic behaviour and that ignoring it overstates closed-loop benefits; and Ponte et al. (2020) demonstrate that the impact of return rates and lead times depends strongly on the degree of information transparency, explicitly noting the "conflicting findings" that pervade the area. A systematic review (Braz et al., 2018) concludes that the primary mitigating factor is an increased return rate and, more strikingly, that *closing* a supply chain can reduce the bullwhip effect—a claim that has lacked a crisp, controlled mechanism.

1.3 Approach: a controlled in-model comparison

This paper answers a deliberately sharp question: *relative to an otherwise identical classical linear supply chain, does a mandatory refurbishment constraint amplify or dampen the bullwhip effect, and through what mechanism?* We answer it with a controlled in-model comparison. Because the linear chain is recovered from the closed-loop model by the single parameter setting $\rho = 0$ (no refurbishment requirement, hence no production constraint), the two systems share an identical demand realisation, ordering policy, lead-time structure and review schedule. Any difference in bullwhip behaviour is therefore attributable solely to the mandatory-refurbishment coupling, not to confounding differences in calibration, demand process or solution method—a level of control that is difficult to achieve when comparing results across separately calibrated studies.

Two further methodological commitments distinguish the analysis. First, every bullwhip ratio is *measured* from the simulated order stream rather than propagated through an analytic variance formula; upstream dispersion therefore responds to the realised dynamics rather than being fixed by construction. Second, we report the *mean-normalised* bullwhip ratio and separate the periodic-review batching component, so that the mechanical amplification caused by bundling several periods of demand into one order pulse is not mistaken for feedback-driven amplification. We compare two ordering policies head-to-head—the classical order-up-to rule and the proportional order-up-to rule (Dejonckheere et al., 2003; Cannella et al., 2021)—over an identical full-factorial design.

1.4 Contributions

1. **A controlled linear-versus-closed-loop comparison.** We compare the classical linear supply chain and a closed-loop chain with a mandatory refurbishment constraint, generated by toggling a single parameter in one model with all other factors held identical, removing the calibration and method confounds that complicate cross-study comparison.

2. **Localised damping and its mechanism.** We show that the constraint leaves downstream echelons statistically unchanged but *dampens* producer bullwhip, that the damping deepens as the constraint binds more tightly, and we identify the truncation of the upper tail of the producer order distribution—a low-pass filtering effect—as the explicit mechanism.

3. **Ordering policy as the primary lever.** Using order-stream measurements rather than analytic propagation, we demonstrate that the ordering policy changes absolute bullwhip levels by an order of magnitude—far more than any single physical parameter—and that the periodic-review interval is the single largest amplifier.

4. **Yield-governed feasibility.** We derive and empirically validate a dimensionally consistent feasibility condition $\phi \cdot \gamma_{ret} \geq 1$ in which refurbishment *yield*, not *delay*, governs whether the constraint binds, and we show the achieved recovered share tracks $\phi\rho$.

5. **A policy-dependent trade-off.** We quantify a tension between variance damping and constraint avoidance: the proportional policy that suppresses the bullwhip effect also thins the refurbished buffer, raising the constraint frequency.

1.5 Structure of the paper

Section 2 positions the work within the bullwhip, control-theoretic and closed-loop literatures and within circular-economy policy. Section 3 develops the model, including the linear chain as a special case, the two ordering policies, the constraint mechanic, the stability condition and the bullwhip metric. Section 4 states the experimental design. Section 5 presents the controlled comparison, the damping mechanism, the role of ordering policy and review interval, the yield-governed feasibility result, and the policy-dependent trade-off. Section 7 discusses implications, reconciles the conflicting signs in the literature, and states the scope and limitations. Section 8 concludes and outlines future work.

2 Related Work and Positioning

2.1 Foundations of the bullwhip effect

The amplification of order variability as one moves upstream was first observed in the system-dynamics tradition of Forrester (1961), whose industrial-dynamics models showed that small fluctuations in retail demand could induce large swings in factory orders through the accumulated effect of feedback delays in ordering and inventory adjustment. The phenomenon was given its enduring managerial articulation by Lee et al. (1997), whose analysis of Procter & Gamble's order patterns for a stable-demand product (disposable diapers) popularised the term "bullwhip" and identified four structural causes: demand-signal processing, in which echelons over-react to demand observations through forecasting; rationing and gaming, in which buyers inflate orders under anticipated shortage; order batching, in which periodic or economic-lot ordering bundles demand into pulses; and price fluctuation, in which promotions induce forward-buying. The four causes remain the canonical taxonomy, and the present model engages two of them directly: demand-signal processing (through the exponential-smoothing forecast) and order batching (through the periodic-review interval).

The behavioural dimension was established by Sterman (1989), whose "beer distribution game"—a tabletop supply-chain simulation played by human participants—demonstrated that decision-makers systematically mis-perceive the feedback delays between ordering and replenishment, under-weighting the supply line of outstanding orders and thereby generating amplification even under stationary demand. Croson and Donohue (2006) later isolated these behavioural causes experimentally, controlling for the operational causes, and quantified the value of inventory information in damping the effect. While our model is operational rather than behavioural—decisions follow a fixed policy rather than human judgement—the Sterman lineage is important context, because it establishes that amplification is not merely a forecasting artefact but a structural property of delayed feedback systems, of which our periodic-review chain is one.

The economic consequences of the bullwhip effect motivate its study: amplified upstream variability inflates safety-stock requirements, induces costly production-rate changes and capacity over-provisioning, degrades service levels through stockouts during demand peaks, and raises the total inventory held across the chain, which Lee et al. (1997) document can reach many weeks of supply in electronics and pharmaceuticals. In a circular-economy setting these costs acquire an environmental dimension as well, since over-production and excess inventory embed wasted material and energy—one

reason the interaction of circularity with the bullwhip effect is of growing interest.

The statistical quantification of the effect is due to Chen et al. (2000), who derived, for a two-stage chain operating an order-up-to policy with exponential-smoothing forecasts, a lower bound on the variance amplification as a function of the smoothing parameter and the lead time, and showed that smoother forecasts reduce—but centralising demand information cannot wholly eliminate—the bullwhip effect. The variance-amplification ratio $\mathrm{Var}(q)/\mathrm{Var}(D)$ introduced in that line of work remains the canonical measure, and our mean-normalised metric (Section 3.7) is its scale-invariant form. Wang and Disney (2016) provide a comprehensive review of the subsequent three decades of progress, tracing the development from these statistical and control-theoretic foundations toward network, behavioural and closed-loop settings. The empirical reality of the effect is more nuanced than the early case evidence suggested: Cachon et al. (2007), examining industry-level US data, find that while bullwhip is pronounced at the firm level, demand can be smoothed rather than amplified once data are aggregated across firms within an industry, a reminder that measured amplification is sensitive to the level of observation and the normalisation used—a methodological point to which our mean-normalised, per-echelon metric responds directly.

2.2 Periodic review and the control-theoretic view

A decisive reframing came from the control-systems tradition. Dejonckheere et al. (2003) modelled replenishment as a feedback controller and proved that the order-up-to rule guarantees a bullwhip effect for any forecasting method, while a *proportional* order-up-to (POUT) rule—which recovers only a fraction $1/T_i$ of the discrepancy between target and actual inventory position each period—can be tuned to damp variability, at the cost of a larger inventory variance. This control lineage traces to the inventory-and-order-based production-control-system (IOBPCS) archetype, whose five-decade development is surveyed by Lin et al. (2017); the proportional feedback structure we adopt was analysed for the order-up-to setting by Disney and Towill (2003), who derived closed-form expressions for the bullwhip and net-stock variance it produces, and by Chen and Disney (2007), who characterised the myopic order-up-to policy under a proportional controller. Disney et al. (2006) make the central trade-off explicit: the same proportional smoothing that tames the bullwhip effect inflates net-stock variance and hence the safety stock required to maintain customer service, so the controller gain encodes a variance-versus-service trade-off rather than a free improvement. The review of replenishment rules and forecasting by Disney and Lambrecht (2008) consolidates these results and provides the variance formulae that our simulation reproduces by direct measurement. Dejonckheere et al. (2004) extend the control-theoretic account to information sharing, showing how much of the amplification survives when downstream demand is made visible upstream. The review interval enters as an order-batching multiplier, consistent with the batching cause of Lee et al. (1997). This policy dichotomy is central to our study: we implement both the OUT and POUT rules exactly and compare them under identical conditions. Gaalman et al. (2022) show, through an eigenvalue analysis of ARMA demand, that the monotonicity of bullwhip in the lead time is itself demand-structure dependent, underlining that quantitative bullwhip results are contingent on modelling assumptions—a caution we adopt explicitly in our scope statement.

2.3 The bullwhip effect in closed-loop supply chains

Closed-loop bullwhip research begins with Tang and Naim (2004) and has grown with the circular-economy agenda, but it has not converged on a consistent sign for the effect of return flows. The field sits within the broader closed-loop supply-chain research programme reviewed by Guide and Van Wassenhove (2009) and Souza (2013), and builds on the hybrid manufacturing/remanufacturing inventory-control foundations of van der Laan et al. (1999), who showed that the choice between push and pull control of the recovery flow materially changes system cost and stock behaviour. Zhou and Disney (2006) establish that returns may either amplify or dampen order and inventory variance depending on the return rate and the remanufacturing lead time, and Zhou et al. (2017) extend this to a multi-echelon chain with remanufacturing uncertainty. Adenso-Díaz et al. (2012) analyse the principal drivers of bullwhip in reverse chains and identify the return rate as a mitigating factor, while Turrisi et al. (2013) show empirically that reverse logistics can be configured to smooth ordering. Dominguez et al. (2020) demonstrate that the *variability* of remanufacturing lead times—rather than their mean—is the damaging quantity, and that models which neglect it overstate the dynamic benefits of circularity.

A unified analytical account is given by Hosoda and Disney (2018), who derive a general theory of closed-loop dynamics and expose a lead-time paradox in which longer remanufacturing delays can, counter-intuitively, reduce order variability; the random-yield and correlation effects underpinning this were established by Hosoda et al. (2015). Of particular relevance to our feasibility analysis, Hosoda et al. (2021) identify a "yield rate paradox" in which the recovery yield governs system behaviour in non-obvious ways—a result our mass-balance condition $\phi\,\gamma_{ret} \geq 1$ complements by making yield the binding feasibility quantity. Ponte et al. (2020) obtain closed-form variance-amplification expressions for four archetypes of information transparency and show that the sign and magnitude of the return-rate and lead-time effects depend on supply-chain visibility, which helps explain the conflicting findings in the literature. Goltsos et al. (2019) synthesise this body of work in a systematic review of remanufacturing dynamics, cataloguing the uncertainties—in returns timing, quantity and quality—that drive closed-loop variability.

Most directly relevant to our modelling choices, Cannella et al. (2021) study proportional order-up-to controllers specifically in closed-loop settings, establishing the POUT family as an appropriate control structure for circular chains; Lin et al. (2022) compare push and pull ordering in a hybrid closed-loop chain using nonlinear (describing-function) analysis; and Dominguez et al. (2019) show that *capacity constraints* dampen closed-loop bullwhip—the mechanism closest to the binding refurbishment constraint we model, differing in that their constraint caps total throughput whereas ours caps only the recovered-content portion. Cannella et al. (2018) analyse capacity restrictions and load-dependent lead times in the same spirit. The mandatory-quota framing itself is anticipated by Ponte et al. (2019), who study the value of *regulating returns* for the dynamic behaviour of hybrid manufacturing–remanufacturing systems, and by Ponte et al. (2022), who isolate the role of batching in closed-loop bullwhip and customer service.

Our contribution within this landscape is methodological as much as substantive: rather than adding another point estimate to a literature of conflicting signs, we hold everything constant except the closed-loop coupling and ask what that coupling alone does. The result—a localised, yield-deepening damping with an explicit distributional mechanism—supplies the missing mechanism behind the review-level ob-

servation (Braz et al., 2018) that closing a chain can reduce the bullwhip effect, and triangulates three near-neighbour studies that each share two of our three defining features but not all three: Cannella et al. (2021) introduce POUT into closed loops but without a binding recovered-content constraint; Dominguez et al. (2019) show capacity constraints dampen closed-loop bullwhip but treat aggregate capacity rather than a refurbishment quota; and Ponte et al. (2019) study regulated returns but not the controlled linear-versus-closed-loop contrast. The combination of a mandatory recovered-content constraint with the classical linear chain recovered as the $\rho = 0$ special case is, to our knowledge, the white space this paper occupies.

2.4 Circular-economy strategies and the policy context

The operational mandate we model is grounded in circular-economy strategy. Potting et al. (2017) organise recovery options into a hierarchy of R-strategies (refuse, rethink, reduce, reuse, repair, refurbish, remanufacture, repurpose, recycle, recover), in which value-retention options higher in the ladder (refurbish, remanufacture) preserve more embedded value than recycling or recovery; Kirchherr et al. (2017) document the conceptual heterogeneity of the field across 114 definitions, and Geissdoerfer et al. (2017) clarify the relationship between the circular economy and sustainability more broadly. Reike et al. (2018) refine the value-retention hierarchy and locate refurbishment precisely among the recovery options, while Morseletto (2020) analyse how quantitative *targets*—the kind of mandated quota our model imposes—are set and operationalised in circular-economy policy. The translation of these concepts into supply-chain practice, including the drivers and barriers firms face, is reviewed by Govindan and Hasanagic (2018) and framed as a supply-chain archetype by Batista et al. (2018); Atasu et al. (2008) establish the strategic and economic rationale that makes refurbishment and remanufacturing attractive to producers independently of mandate. Refurbishment, the recovery mode in our model, is sometimes termed "light remanufacturing": it restores a used product to a specified working and cosmetic standard, typically through the replacement of parts rather than full disassembly. Mandatory quotas of the kind we model arise directly from EPR instruments such as the WEEE Directive, whose binding recovery targets transfer end-of-life responsibility—and, increasingly, recovered-material sourcing obligations—onto producers. The empirical backdrop underscores why the binding case matters: Baldé et al. (2024) report that only about 22% of global electronic waste was documented as formally collected and recycled in 2022, far below the targets set by instruments such as the WEEE Directive. This persistent shortfall of actual recovery relative to target makes the binding-constraint regime studied here the empirically relevant case rather than a theoretical edge case.

3 Model

This section develops the discrete-event Monte-Carlo model in full. The notation is collected in Table 1. We describe the physical structure (Section 3.1), the demand and forecasting processes (Section 3.2), the inventory policies (Section 3.3), the refurbishment constraint and the recovery of the linear chain as a special case (Section 3.4), the steady-state feasibility condition (Section 3.5), and the bullwhip metric and its normalisation (Section 3.7). The sequence of events within a period is summarised at the end (Section 3.8).

Table 1: Notation.

Symbol	Meaning
$i \in \{R, W, P\}$	echelon: retailer, wholesaler, producer
D_t	customer demand in period t
μ_D, σ_D, CV_D	mean, standard deviation, coefficient of variation of demand
t_i	review interval of echelon i (factor t for W and P; t_R=1)
L_i	physical lead time of echelon i (L_R=1, L_W=2, L_P=2)
α	exponential smoothing parameter
$\hat{\mu}_{i,t}, \hat{\sigma}_{i,t}$	smoothed mean and standard deviation of observed signal
z	service factor (1.645, 95%)
$S_{i,t}$	order-up-to target
$IP_{i,t}$	inventory position (on-hand + pipeline − backlog)
$O_{i,t}$	order placed by echelon i
T_i	proportional controller horizon (POUT)
ρ	return / mandated recovered-content rate
τ	refurbishment processing delay (periods)
ϕ	refurbishment yield (success rate)
I_t^{ref}	available recovered-component stock
Q_t^{act}	actual production after applying the constraint
BWE_i	mean-normalised bullwhip ratio at echelon i
$\beta_P(t)$	periodic-review batching factor
AR_i	amplification ratio, CLSC relative to linear
γ_{ret}	effective recovered coverage ratio

3.1 Physical structure

We model a serial four-echelon chain consisting of a Retailer (R), a Wholesaler (W) and a Producer (P) in the forward direction, supported by a Refurbishment station that processes customer returns and feeds recovered components back to the producer. Customer demand D_t is the only exogenous driver of the system. Material flows downstream from producer to customer; order information flows upstream from customer to producer; and returned units flow from the customer into the refurbishment station and re-enter production as recovered components. Each forward echelon $i \in \{R, W, P\}$ holds inventory, observes the order stream of its immediate downstream partner (the retailer observes customer demand), forecasts, and places replenishment orders on a periodic-review schedule with review interval t_i and physical lead time L_i.

We use echelon-specific lead times $L_R = 1$, $L_W = 2$, $L_P = 2$ (periods), reflecting the typical shortening of replenishment lead times toward the customer end of the chain, and a retailer review interval $t_R = 1$. The wholesaler and producer share a common review interval t, which is a design factor. All quantities are expressed in abstract demand units; the model is calibration-light by design, since the controlled comparison depends on relative rather than absolute magnitudes.

3.2 Demand and forecasting

Customer demand is independent and identically distributed across periods, normally distributed and truncated at zero,

$$D_t \sim \max\{0,\, \mathcal{N}(\mu_D,\, \sigma_D^2)\}, \qquad \mu_D = 100, \quad \sigma_D = CV_D\, \mu_D, \quad CV_D = 0.20. \tag{1}$$

The truncation, implemented by resampling the rare negative draws, preserves non-negativity without materially distorting the moments at $CV_D = 0.20$. A fixed random seed per replication makes each run exactly reproducible and ensures that the linear and closed-loop systems experience an identical demand realisation within a replication, which is what makes the comparison controlled at the level of individual sample paths.

Each echelon maintains exponentially smoothed estimates of the mean and variance of the signal it observes ($y_{i,t}$, equal to customer demand for the retailer and to the downstream order for the wholesaler and producer):

$$\hat{\mu}_{i,t} = \alpha\, y_{i,t} + (1-\alpha)\, \hat{\mu}_{i,t-1}, \tag{2}$$

$$\hat{\sigma}^2_{i,t} = \alpha\, \big(y_{i,t} - \hat{\mu}_{i,t}\big)^2 + (1-\alpha)\, \hat{\sigma}^2_{i,t-1}, \tag{3}$$

with smoothing parameter $\alpha \in (0,1)$ common to all echelons. The decisive methodological point is that $\hat{\sigma}_{i,t}$ in Eq. (3) is estimated from the *realised* order stream that the echelon actually receives. It is therefore an emergent quantity: as amplification builds upstream, each echelon observes and forecasts a more volatile signal, and its safety stock responds accordingly. This contrasts with variance-propagation models in which the upstream standard deviation is computed analytically from end-customer demand and a fixed amplification factor, which forces the dispersion to be a function of the demand process alone and cannot represent the feedback through which the ordering policy and the constraint shape the variance actually seen upstream.

3.3 Inventory policy

Each echelon targets an order-up-to level covering the risk period of review interval plus lead time:

$$S_{i,t} = \hat{\mu}_{i,t}\, (t_i + L_i) + z\, \hat{\sigma}_{i,t}\, \sqrt{t_i + L_i}, \qquad z = 1.645 \text{ (95\% service factor)}. \tag{4}$$

The first term covers expected demand over the risk period; the second is the safety stock, scaled by the service factor z and the square-root-of-risk-period rule standard in periodic-review theory. The inventory position is

$$IP_{i,t} = I_{i,t} + \textstyle\sum_{\text{pipeline}} - B_{i,t}, \tag{5}$$

that is, on-hand inventory plus in-transit (pipeline) receipts minus backlog. On each review epoch ($t \bmod t_i = 0$) the echelon places an order; between review epochs it places none. We study two replenishment rules.

Classical order-up-to (OUT). The order closes the entire gap between target and inventory position:

$$O_{i,t}^{\text{OUT}} = \max\big(S_{i,t} - IP_{i,t},\ 0\big). \tag{6}$$

This is the rule analysed by Chen et al. (2000) and shown by Dejonckheere et al. (2003) to guarantee amplification.

Proportional order-up-to (POUT). The order recovers only a fraction $1/T_i$ of the inventory gap, added to the demand forecast:

$$O_{i,t}^{\text{POUT}} = \max\Big(\hat{\mu}_{i,t} + \tfrac{1}{T_i}\big(S_{i,t} - IP_{i,t}\big),\ 0\Big), \tag{7}$$

with proportional controller horizon $T_i \geq 1$. As $T_i \to 1$ the POUT rule reduces to OUT; for $T_i > 1$ it smooths the order stream, trading lower order variance for higher inventory variance. We use $T_i = 4$ throughout, a moderate value well inside the stable region of the proportional controller (Dejonckheere et al., 2003; Cannella et al., 2021). The proportional rule is the appropriate control structure for closed-loop chains established by Cannella et al. (2021), and its inclusion lets us separate effects that are intrinsic to the closed loop from effects that are artefacts of an aggressive ordering rule.

3.4 Refurbishment constraint and the linear chain

Customer returns are generated in proportion to units sold and enter the refurbishment station at rate ρ (the return/mandate rate). They emerge after a processing delay of τ periods and become usable recovered components with yield (success rate) $\phi \in (0, 1]$; a fraction $1 - \phi$ is lost to inspection failure or irreparable damage. The producer must source a fraction ρ of each production release from recovered components, reflecting a mandated recovered-content quota. Writing $O_{P,t}^{\text{rel}}$ for the producer's release decision (from Eq. (6) or (7)), the constraint mechanic is

$$r_t^{\text{need}} = \rho\, O_{P,t}^{\text{rel}}, \qquad n_t^{\text{need}} = (1-\rho)\, O_{P,t}^{\text{rel}}, \tag{8}$$

$$r_t^{\text{used}} = \min\big(I_t^{\text{ref}},\ r_t^{\text{need}}\big), \qquad Q_t^{\text{act}} = r_t^{\text{used}} + n_t^{\text{need}}, \tag{9}$$

where I_t^{ref} is the available recovered-component stock. The new-build portion n_t^{need} is always produced; only the recovered portion can fall short. A period is *constrained* when $r_t^{\text{used}} < r_t^{\text{need}}$, i.e. when the recovered stock cannot meet the mandated recovered requirement and actual production Q_t^{act} falls below the desired release. This binding-from-below mechanic mirrors the load-dependent capacity restriction of Cannella et al. (2018), specialised to a recovered-content quota.

Recovery of the linear chain. Setting $\rho = 0$ removes the recovered requirement entirely: $r_t^{\text{need}} = 0$, the constraint can never bind, and $Q_t^{\text{act}} = O_{P,t}^{\text{rel}}$. The model then reduces exactly to a standard four-echelon periodic-review forward supply chain. This is the linear benchmark used throughout. Crucially, the benchmark and the closed-loop system are the same code under different ρ, so they share the demand process, forecasting, ordering rule, lead times and review schedule; the comparison is free of calibration or methodological confounds. We denote the closed-loop system by CLSC and reserve "linear" for the

$\rho = 0$ case.

3.5 Steady-state feasibility condition

For the closed loop to operate without persistently binding constraints, the recovered inflow must, in the mean, cover the mandated recovered requirement. Over the long run, returns enter at rate ρ of sales and become usable at yield ϕ, so the mean usable recovered inflow is $\phi\,\rho\,\bar{D}$ (with $\bar{D}$ the mean throughput, which in steady state equals mean demand). The mandated recovered draw is $\rho\,\bar{Q}_P$, and since the bullwhip effect inflates the *variance* of orders but not their *mean*, mean production satisfies $\bar{Q}_P \approx \bar{D}$. Feasibility therefore requires

$$\phi\,\rho\,\bar{D} \;\geq\; \rho\,\bar{Q}_P \quad\Longleftrightarrow\quad \phi\,\gamma_{ret} \;\geq\; 1, \tag{10}$$

where $\gamma_{ret} = \rho\bar{D}/(\rho\bar{Q}_P) \to 1$ is the effective recovered coverage ratio. The condition $\phi\,\gamma_{ret} \geq 1$ is dimensionally consistent (a yield times a coverage ratio, both dimensionless, compared with unity) and makes clear that feasibility is governed by yield. When $\phi < 1$ the recovered inflow structurally undershoots the mandate and the constraint binds with a frequency that grows as ϕ falls; the achieved recovered share of production is then $\phi\rho$ rather than the target ρ. The bullwhip effect enters Eq. (10) only indirectly, through the size of the recovered buffer needed to absorb order-stream dispersion; it is not itself a feasibility threshold. We validate this condition empirically in Section 5.5.

3.6 Why variance is measured, not propagated

A recurring methodological fork in the bullwhip literature is whether the upstream order variance is *measured* from the realised order stream or *propagated* analytically from end-customer demand through a closed-form amplification factor. The distinction is not cosmetic, and it matters especially in a closed-loop setting, so we state our choice and its consequences explicitly.

In the propagation approach, the standard deviation seen by an upstream echelon is written as the demand standard deviation multiplied by a lead-time- and review-dependent amplification factor derived under an order-up-to policy with a specified forecasting rule. This yields elegant closed forms and is exact under its own assumptions. But it has a structural consequence: the upstream dispersion becomes a fixed function of the demand process and the lead-time parameters alone. It cannot respond to anything that changes the order stream without changing those inputs—including a change in the ordering rule, or the introduction of a production constraint that clips the order distribution. In a closed-loop model that propagates variance analytically, a recovered-content constraint that reduces order volume will mechanically reduce the propagated dispersion as well, because the safety-stock term is tied to a demand-derived standard deviation; the "damping" observed would then be partly an artefact of the propagation assumption rather than an emergent property of the dynamics.

Our model takes the measurement approach: each echelon estimates the standard deviation of the signal it actually observes (Eq. 3), and the bullwhip ratio is computed from the realised order series after the simulation has run. Consequently, the dispersion at each echelon is an outcome, not an input. When the ordering rule changes from OUT to POUT, or when the constraint clips the producer order tail, the measured variance reflects those changes directly, and the amplification ratio AR_i isolates the effect of

the closed-loop coupling without contamination from a built-in dispersion assumption. This is the reason the controlled comparison can be trusted to attribute the producer damping to the constraint mechanism rather than to a modelling convention: the only difference between the linear and closed-loop runs is the value of ρ, and the variance in both is measured the same way from the same kind of order stream.

3.7 Bullwhip metric and normalisation

We quantify amplification with the *mean-normalised* bullwhip ratio, measured per echelon over the post-warm-up window,

$$BWE_i \;=\; \frac{\mathrm{Var}(O_i)\,/\,\mathrm{E}[O_i]^2}{\mathrm{Var}(D)\,/\,\mathrm{E}[D]^2} \;=\; \frac{CV^2(O_i)}{CV^2(D)}, \tag{11}$$

i.e. the ratio of the squared coefficient of variation of orders to that of demand. Normalisation by the squared mean is essential for two reasons. First, it makes the measure scale-invariant and therefore comparable across echelons whose order magnitudes differ. Second, and more importantly in a periodic-review setting, the raw variance ratio $\mathrm{Var}(O)/\mathrm{Var}(D)$ is inflated whenever orders are batched: an echelon reviewing every t periods bundles roughly t periods of demand into a single order pulse, multiplying the raw order variance by a factor of order t even in the complete absence of feedback amplification. Reporting the un-normalised ratio therefore confounds batching with genuine amplification and produces ratios that can be orders of magnitude larger than the single-digit values expected for serial chains. Eq. (11) is the scale-invariant form of the canonical $\mathrm{Var}(q)/\mathrm{Var}(D)$ measure of Chen et al. (2000).

To isolate the batching contribution explicitly we report the *batching factor*

$$\beta_P(t) \;=\; \frac{BWE_P(t)}{BWE_P(1)}, \tag{12}$$

the producer bullwhip at review interval t relative to a continuous-equivalent reference ($t = 1$). To quantify the closed-loop effect against the linear benchmark we report the *amplification ratio*

$$AR_i \;=\; \frac{BWE_i^{\text{CLSC}}}{BWE_i^{\text{linear}}}, \tag{13}$$

where $AR_i < 1$ indicates that the closed-loop constraint dampens variability at echelon i, $AR_i = 1$ indicates no effect, and $AR_i > 1$ indicates amplification. The amplification ratio is the central quantity of the controlled comparison because it cancels everything the two systems share and exposes only the effect of the closed-loop coupling.

3.8 Sequence of events

Within each period the model executes the following steps, consistent with the order-up-to event sequence of Chen et al. (2000):

1. Customer demand D_t is realised; the retailer updates its forecast (Eqs. 2–3), receives the pipeline arrival due this period, serves demand (backlogging any shortfall), and on a review epoch places an order $O_{R,t}$ (Eq. 6 or 7).
2. The wholesaler observes $O_{R,t}$, updates its forecast, receives its pipeline arrival, ships to the retailer

Table 2: Full-factorial experimental design (per policy). The linear benchmark sets $\rho = 0$; all other factors and fixed parameters are identical between the linear and closed-loop systems.

Factor	Levels	Count
Smoothing parameter α	$\{0.2, 0.4, 0.6, 0.8\}$	4
Review interval t (W and P)	$\{1, 2, 3\}$	3
Return / mandate rate ρ	$\{0.2, 0.3, 0.4, 0.5\}$	4
Refurbishment delay τ	$\{1, 3, 5, 10\}$	4
Success (yield) rate ϕ	$\{0.6, 0.8, 1.0\}$	3
Configurations		576
Replications per configuration		30
Runs per policy		17,280
Total runs (two policies)		34,560

(the shipment becomes the retailer's pipeline receipt after L_R), and on a review epoch places an order $O_{W,t}$.

3. The producer observes $O_{W,t}$, updates its forecast, and receives its work-in-process arrival. The refurbishment pipeline advances: returns deposited τ periods earlier emerge and are added to the recovered stock at yield ϕ; new returns, proportional to current sales, enter the pipeline.

4. The producer ships finished goods to the wholesaler from on-hand inventory. On a review epoch it computes a desired release (Eq. 6 or 7), applies the refurbishment constraint (Eqs. 8–9), and the actual build Q_t^{act} enters the work-in-process pipeline, arriving after L_P.

5. All state variables (orders, inventories, recovered stock, constraint indicator, production split) are recorded.

Warm-up periods are discarded before any metric is computed.

4 Experimental Design

We run a full-factorial design over both ordering policies, summarised in Table 2. The five factors—smoothing parameter α, review interval t, return/mandate rate ρ, refurbishment delay τ, and success rate ϕ—span $4 \times 3 \times 4 \times 4 \times 3 = 576$ configurations. Each configuration is replicated 30 times with independent demand seeds, giving $576 \times 30 = 17{,}280$ runs per policy and $34{,}560$ runs in total. Each run spans $T = 500$ periods with a 100-period warm-up; steady-state metrics are computed over periods 101–500. Point estimates are replication means, and we report 95% confidence intervals obtained by bootstrap resampling (2,000 resamples) over the 30 replications. The linear benchmark is generated from the same model and design with $\rho = 0$.

The fixed parameters, common to all runs, are: mean demand $\mu_D = 100$ and coefficient of variation $CV_D = 0.20$; echelon lead times $L_R = 1$, $L_W = 2$, $L_P = 2$; retailer review interval $t_R = 1$; service factor $z = 1.645$; and proportional controller horizon $T_i = 4$. We define a *balanced operating point* ($\alpha = 0.4$, $t = 2$, $\rho = 0.3$, $\tau = 3$, $\phi = 0.8$) that we use as the reference configuration when illustrating per-echelon results; it lies in the interior of the design and exhibits a moderate, non-trivial constraint frequency.

4.1 A worked single-period illustration

To make the mechanics concrete, consider a single review period at the producer under the balanced operating point, with the OUT rule. Suppose the producer's smoothed mean is $\hat{\mu}_P = 100$ and smoothed standard deviation $\hat{\sigma}_P = 28$ (the latter inflated above the demand $\sigma_D = 20$ by the upstream amplification it has accumulated), with review interval $t = 2$ and lead time $L_P = 2$, so the risk period is $t + L_P = 4$. The order-up-to target is

$$S_P = \hat{\mu}_P(t + L_P) + z\,\hat{\sigma}_P\sqrt{t + L_P} = 100 \cdot 4 + 1.645 \cdot 28 \cdot 2 = 400 + 92.1 = 492.1.$$

If the inventory position has been drawn down to $IP_P = 300$ by a run of high demand, the desired release is $O_P^{\text{rel}} = \max(492.1 - 300, 0) = 192.1$ units. In the linear chain this is produced in full. In the closed-loop chain with $\rho = 0.3$, the mandated recovered draw is $r^{\text{need}} = 0.3 \cdot 192.1 = 57.6$ units and the new-build portion is $n^{\text{need}} = 134.5$ units. If the recovered stock currently holds only $I^{\text{ref}} = 40$ units, then $r^{\text{used}} = \min(40, 57.6) = 40$, the period is constrained, and actual production is $Q^{\text{act}} = 40 + 134.5 = 174.5$ units—a release of 174.5 rather than the desired 192.1. The 17.6-unit shortfall is exactly the top of the desired release that could not be sourced; the large upward excursion is clipped. Repeated over many periods, it is this clipping of the largest releases—and only the largest, since small releases are never recovered-constrained—that thins the upper tail of the producer order distribution and lowers the measured bullwhip ratio, while leaving the mean throughput essentially unchanged because the new-build portion always flows.

5 Results

We present the results in the order of the paper's logic. Section 5.1 reports the controlled linear-versus-closed-loop comparison, the central result. Section 5.2 establishes the damping mechanism through the order distribution and trajectory. Section 5.3 examines the absolute levels of bullwhip and the roles of ordering policy and review interval. Section 5.5 validates the yield-governed feasibility condition. Section 5.6 reports the policy-dependent trade-off. Throughout, "balanced operating point" refers to the reference configuration defined in Section 4.

5.1 The controlled comparison: localised damping

Figure 1 contrasts echelon bullwhip between the linear chain ($\rho = 0$) and the closed-loop chain ($\rho = 0.3$) at the balanced operating point, holding everything else identical. The pattern is unambiguous and is the same under both ordering policies. At the retailer and the wholesaler the two systems are statistically indistinguishable: the confidence intervals overlap almost completely, because the refurbishment constraint acts only at the producer and therefore cannot alter the order streams generated downstream of it. At the producer, by contrast, the closed-loop chain exhibits *lower* bullwhip than the linear chain: producer BWE falls from 198.2 to 168.7 under OUT and from 31.5 to 29.2 under POUT. The summary figures and their confidence intervals are collected in Table 3.

The per-echelon amplification ratio (Eq. 13) makes the comparison precise. Figure 2 shows $AR_R \approx AR_W \approx 1.00$ for both policies, confirming the absence of any closed-loop effect downstream, while

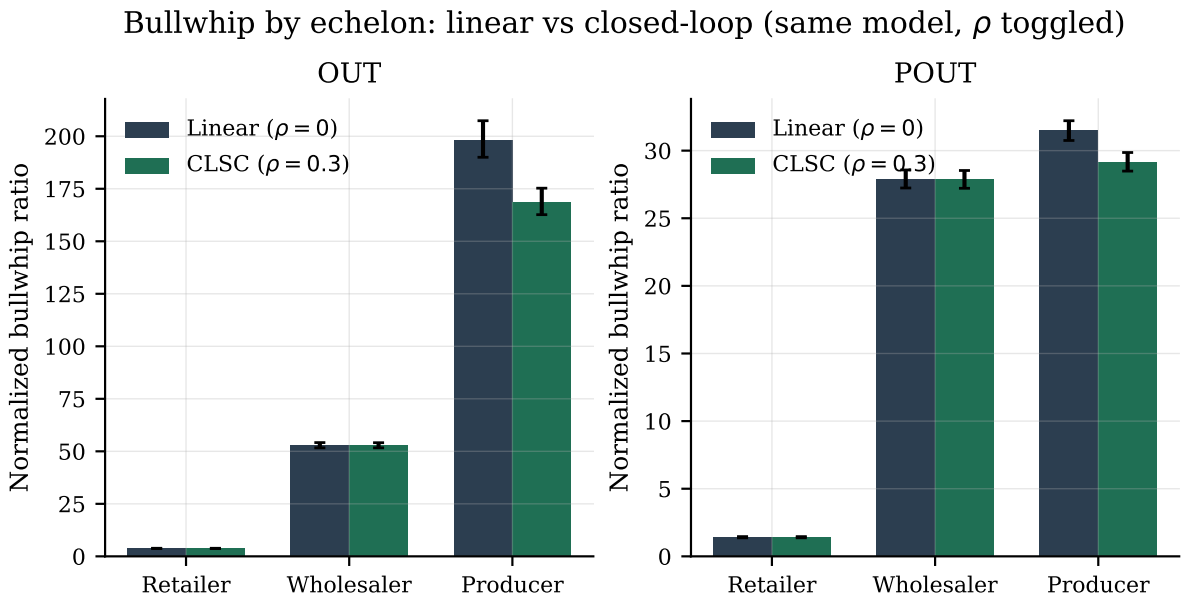


Figure 1: Bullwhip by echelon, linear (ρ=0) versus closed-loop (ρ=0.3), generated from the same model with all other factors held identical, at the balanced operating point. Downstream echelons are unaffected; the producer is damped under the closed-loop constraint. Error bars are 95% bootstrap confidence intervals.

Table 3: Headline results at the balanced operating point (α=0.4, t=2, ρ=0.3, τ=3, ϕ=0.8). BWE values are mean-normalised bullwhip ratios (Eq. 11); the linear system sets ρ=0. Constraint rate is the fraction of periods in which the producer is bindingly constrained.

Policy	System	BWE_R	BWE_W	BWE_P	Constraint rate
OUT	Linear	3.77	52.85	198.22	0.000
	CLSC	3.77	52.87	168.67	0.074
POUT	Linear	1.41	27.87	31.50	0.000
	CLSC	1.41	27.87	29.17	0.366

$AR_P = 0.85$ under OUT and 0.93 under POUT. The mandatory refurbishment constraint is therefore not a source of additional bullwhip relative to the linear chain at this operating point; if anything it is a mild stabiliser of producer ordering. This is the first principal finding: *the closed-loop coupling does not amplify the bullwhip effect and is localised entirely at the constrained echelon.*

That the effect is confined to the producer deserves emphasis, because it runs against a common intuition that disturbances introduced upstream should propagate and be amplified downstream as well. Here the causal direction is the reverse of the bullwhip cascade: order information flows from customer to producer, so a modification of the producer's ordering behaviour has no downstream order stream to contaminate. The wholesaler and retailer continue to face exactly the same inputs they faced in the linear chain. The closed loop modifies only the most upstream node, and it does so by limiting that node's own production releases.

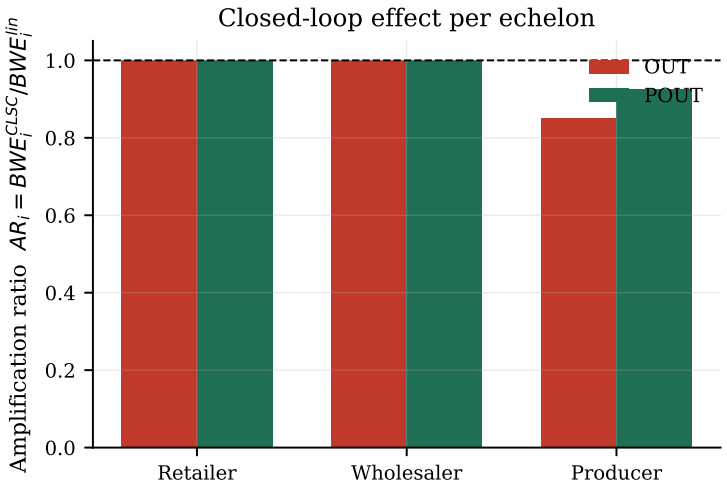


Figure 2: Amplification ratio $AR_i = BWE_i^{\text{CLSC}}/BWE_i^{\text{linear}}$ per echelon, both policies. Values below the dashed unity line indicate damping; the effect is confined to the producer.

5.2 The mechanism: a low-pass filter on producer orders

Why should a binding constraint *reduce* measured variability rather than increase it? The answer is distributional. Figure 3 shows representative producer-order trajectories over a 100-period window for both policies, overlaying the customer demand, the linear-chain producer order, and the closed-loop producer order at a low yield ($\phi = 0.7$) that makes the constraint bind frequently. The linear-chain order series exhibits large upward spikes—the catch-up pulses that the order-up-to logic issues to refill inventory after a run of high demand. In the closed-loop series these upward excursions are visibly clipped: when a large release is desired but recovered stock is insufficient, production is capped and the spike is truncated.

The effect on the order distribution is shown directly in Figure 4, which plots the density of positive producer orders for the linear and closed-loop systems. Under both policies the closed-loop distribution is concentrated more tightly around its centre and carries appreciably less mass in the upper tail: the constraint removes precisely the large orders, leaving the bulk of the distribution intact. Because variance is dominated by the tails, removing upper-tail mass lowers the variance and hence the normalised bullwhip ratio. The constraint thus behaves as a *low-pass filter* on the producer's order signal: it attenuates the high-amplitude excursions while passing the low-frequency bulk. This is a filtering effect, not a control action—the producer does not choose to smooth its orders; rather, the physical unavailability of recovered components prevents the largest releases from being executed.

This mechanism explains the yield dependence of the damping reported next. The more often the constraint binds—i.e. the lower the yield ϕ—the more of the upper tail is removed, and the deeper the damping. Figure 5 traces the producer amplification ratio AR_P and the constraint frequency as ϕ varies from 0.6 to 1.0. As ϕ falls and the constraint binds more often, AR_P declines monotonically below unity, reaching 0.77 under OUT and 0.91 under POUT at $\phi = 0.6$. At $\phi = 1.0$ the constraint essentially never binds (it is feasible by Eq. 10) and $AR_P \to 1$, recovering the linear chain as a limiting case—an internal consistency check that the model passes exactly.

Figure 6 generalises the yield dependence across the (ϕ, ρ) plane. The damping deepens (greener

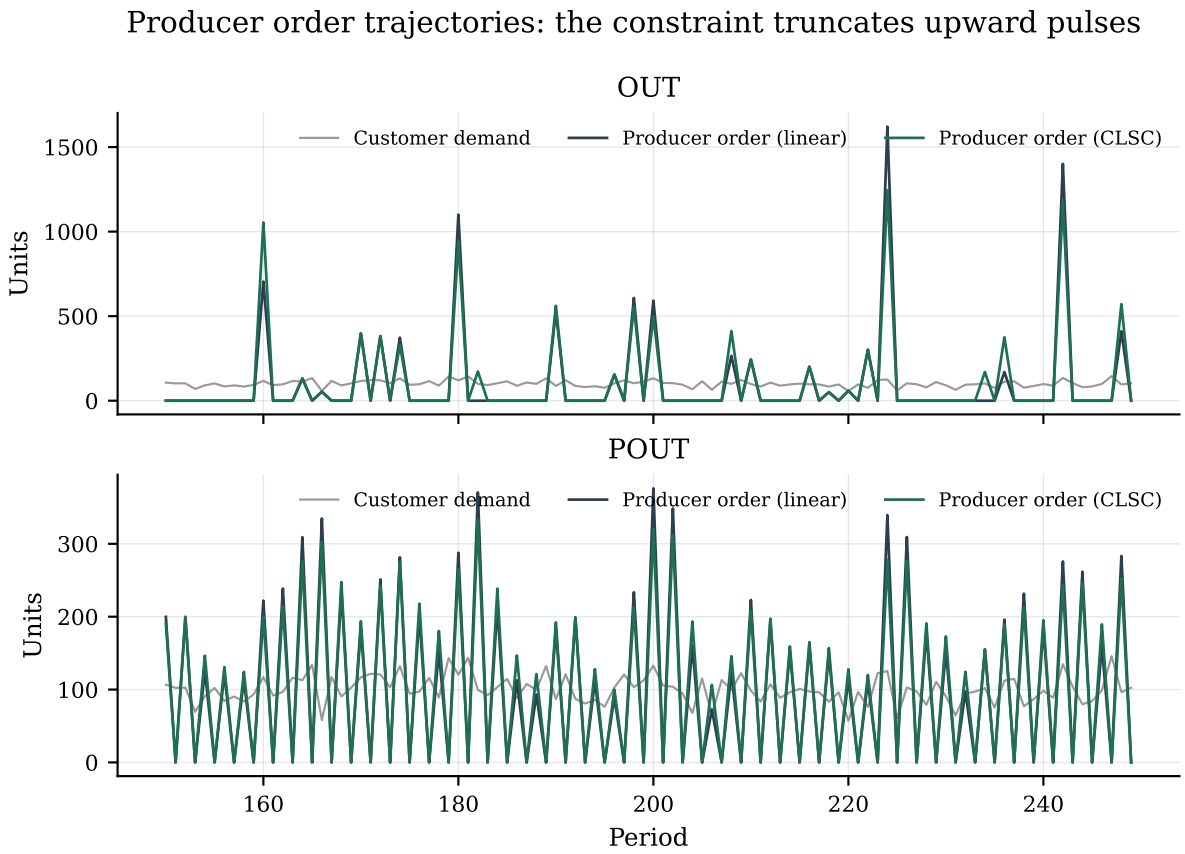


Figure 3: Producer order trajectories over a representative window (periods 150–250), customer demand versus linear and closed-loop producer orders, for each policy (α=0.4, t=2, ρ=0.3, τ=3; closed loop at ϕ=0.7). The closed-loop constraint clips the upward catch-up pulses.

cells, lower AR_P) as yield falls and as the mandate rate rises, both of which increase the frequency and severity of binding. The surface is smooth and monotone, consistent with the single distributional mechanism identified above; there is no region in which the closed-loop coupling amplifies producer bullwhip ($AR_P > 1$) at this operating slice.

5.3 Absolute levels: ordering policy and review interval

While the closed-loop effect on bullwhip is modest, the *absolute* level of bullwhip is governed primarily by the ordering policy. Across the design the switch from OUT to POUT reduces echelon bullwhip by roughly an order of magnitude. At the balanced operating point (Table 3) the producer ratio falls from $\sim$169 under OUT to $\sim$29 under POUT, and the retailer ratio from 3.8 to 1.4. Figure 7 shows the propagation across echelons under each policy: the qualitative cascade—increasing amplification from retailer to producer—is identical in shape, but POUT holds the entire cascade at a far lower level. The variability cascade in coefficient-of-variation terms is shown in Figure 8, which makes the same point on the scale-invariant CV axis: POUT compresses the upstream growth of variability.

The periodic-review interval is the single largest amplifier of producer variability. Figure 9 reports the batching factor $\beta_P(t)$ (Eq. 12) in a near-linear configuration. Relative to the continuous-equivalent $t = 1$ reference, the review interval multiplies producer bullwhip by 1.6$\times$ at $t = 2$ and 2.2$\times$ at $t = 3$

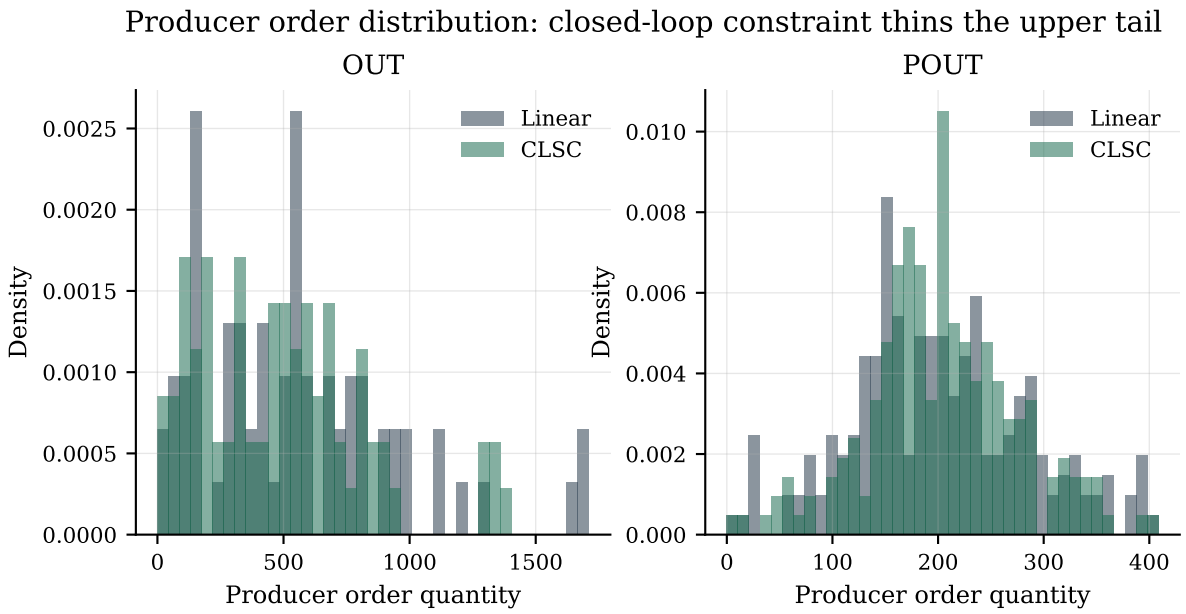


Figure 4: Producer order distribution (positive orders), linear versus closed-loop, each policy. The closed-loop constraint thins the upper tail, which is what lowers the measured order variance.

Table 4: Producer bullwhip BWE_P over the joint grid of review interval t and smoothing parameter α, averaged over all other factors, for each policy.

		Smoothing parameter α			
Policy	t	0.2	0.4	0.6	0.8
OUT	1	19.7	101.2	172.8	216.0
	2	59.2	166.6	260.3	373.3
	3	102.1	231.1	355.6	538.0
POUT	1	1.0	2.9	6.3	12.6
	2	27.0	29.5	34.2	43.1
	3	54.6	57.2	61.1	69.9

under OUT, and by $8.2\times$ and $15.7\times$ under POUT. The much larger multiplier under POUT is expected: because POUT's $t = 1$ baseline producer bullwhip is small (a realistic 3.9), the additive variance injected by batching represents a larger relative increase. In absolute terms the batched POUT producer bullwhip remains far below the OUT level. The practical message is that a large part of any apparent "amplification" observed in periodic-review chains is a mechanical consequence of bundling demand into batched order pulses, and is removed neither by closing the loop nor by choosing a particular yield; it is removed only by reviewing more frequently or by reporting a batching-corrected metric.

5.4 Interaction effects across the full design

The marginal effects of the factors (Appendix B) do not by themselves reveal how the factors combine. Table 7 reports the producer bullwhip over the joint (t, α) grid—the two factors with the largest main effects—for each policy, averaged over the remaining factors. Three features of the interaction are worth drawing out.

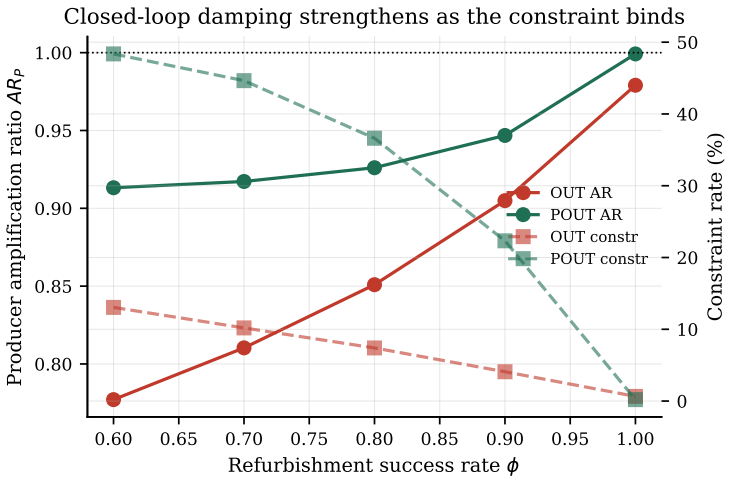


Figure 5: Producer amplification ratio AR_P (solid) and constraint frequency (dashed) versus refurbishment yield ϕ. Lower yield means the constraint binds more often, removes more upper-tail order mass, and deepens the damping. At ϕ=1 the linear chain is recovered.

First, the two amplifying factors compound rather than merely add: under OUT, the worst corner ($t = 3, \alpha = 0.8$) reaches $BWE_P = 538$, far above what a sum of the separate t and α effects would predict, because an aggressive forecast generates large gaps and a long review interval batches them into still larger pulses. The bullwhip effect is therefore most dangerous when responsive forecasting and infrequent review coincide. Second, the proportional policy changes the *structure* of the interaction, not only its level. Under POUT at $t = 1$ the producer bullwhip is close to unity (1.0–12.6 across α), meaning that a frequently reviewed proportional chain barely amplifies at all, even with an aggressive forecast; the α sensitivity that dominates the OUT column is almost entirely suppressed. Third, once batching is present ($t \geq 2$), the POUT advantage narrows in relative terms because the batching variance is common to both policies—confirming the earlier observation that the review interval is a policy-independent amplifier. The managerial reading is that the proportional policy and frequent review are complementary: the largest gains come from adopting both, since POUT's smoothing is most effective precisely when it is not overwhelmed by batching.

5.5 Feasibility is governed by yield, not delay

Figure 10 reports the constraint frequency as a function of the mandate rate ρ (left) and the refurbishment delay τ (right). The constraint frequency rises steadily with ρ, as a larger mandated recovered content is harder to satisfy from a fixed recovered inflow. It is, by contrast, comparatively insensitive to the refurbishment delay τ: lengthening the processing delay shifts the timing of recovered arrivals but, in steady state, does not change their mean rate, so it has only a weak second-order effect on how often the constraint binds. The decisive lever is the yield ϕ. As established in Figure 5, the constraint frequency collapses to near zero at $\phi = 1.0$ and rises sharply as ϕ falls; correspondingly, the achieved recovered share of production equals $\phi\rho$ (reaching the target $\rho = 0.30$ at $\phi = 1.0$, falling to 0.24 at $\phi = 0.8$ and 0.18 at $\phi = 0.6$). This is exactly the prediction of the mass-balance condition Eq. (10): feasibility is a yield phenomenon. The empirical collapse of the constraint at $\phi = 1$ and its severity at low ϕ, together with the $\phi\rho$ scaling of the achieved share, constitute a direct validation of the condition.

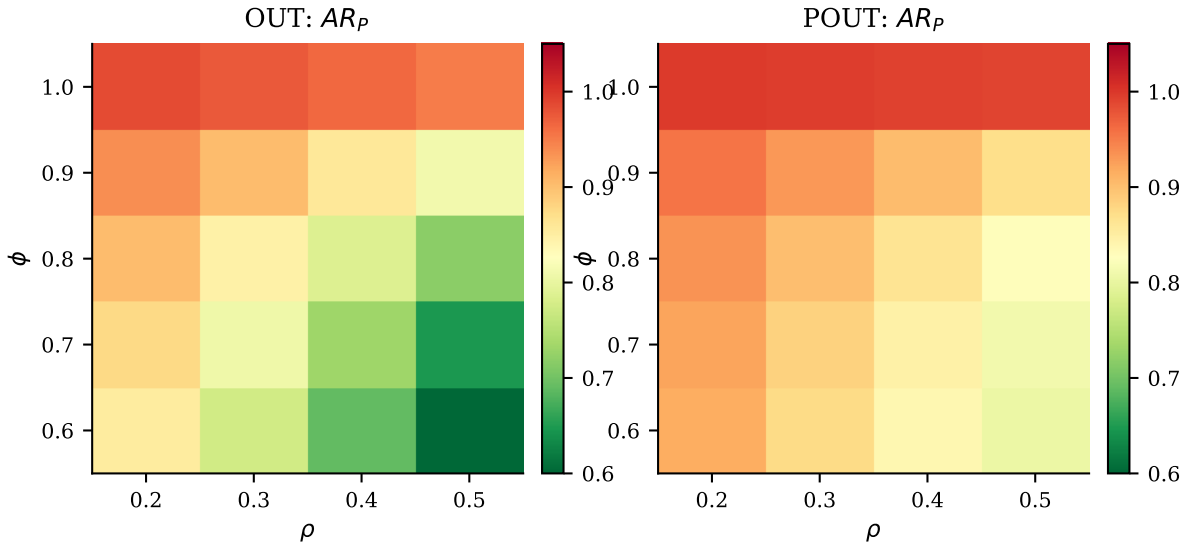


Figure 6: Producer amplification ratio AR_P over the (ϕ, ρ) plane for each policy. Damping deepens as the constraint binds more tightly (lower ϕ, higher ρ). No region shows amplification.

5.6 A policy-dependent trade-off

The damping benefit of the proportional policy carries an operational cost. Because POUT issues smoother and smaller orders, it accumulates less recovered buffer between binding episodes, and it therefore operates at a substantially *higher* constraint frequency than OUT: averaged over the design, the constraint binds in 31% of periods under POUT versus 10% under OUT, and at the balanced operating point the figures are 36.6% versus 7.4% (Table 3). Variance damping and constraint avoidance are thus partially in tension. The very mechanism that suppresses the bullwhip effect—smoother, smaller orders—also thins the recovered buffer that protects against binding the recovered-content constraint. A firm that adopts a proportional controller to stabilise its order stream will, other things equal, find its mandated recovered-content quota harder to meet period by period, and must compensate with additional recovered-stock buffering or a higher yield.

Figure 11 maps producer bullwhip over the (ρ, τ) plane for both policies and shows that the low-bullwhip region is far larger under POUT; and Figure 12 plots producer bullwhip against the achieved recovered share, exhibiting the trade-off frontier traced out across the design. Together they show that the policy choice is not a free lunch: it relocates the operating point along a frontier that trades order stability against constraint exposure.

5.7 Robustness and internal validation

Before summarising, we report the internal-consistency checks that establish the model behaves as its construction requires. These are not parameter studies but verifications that the simulation conserves the quantities it should and reproduces known limiting cases.

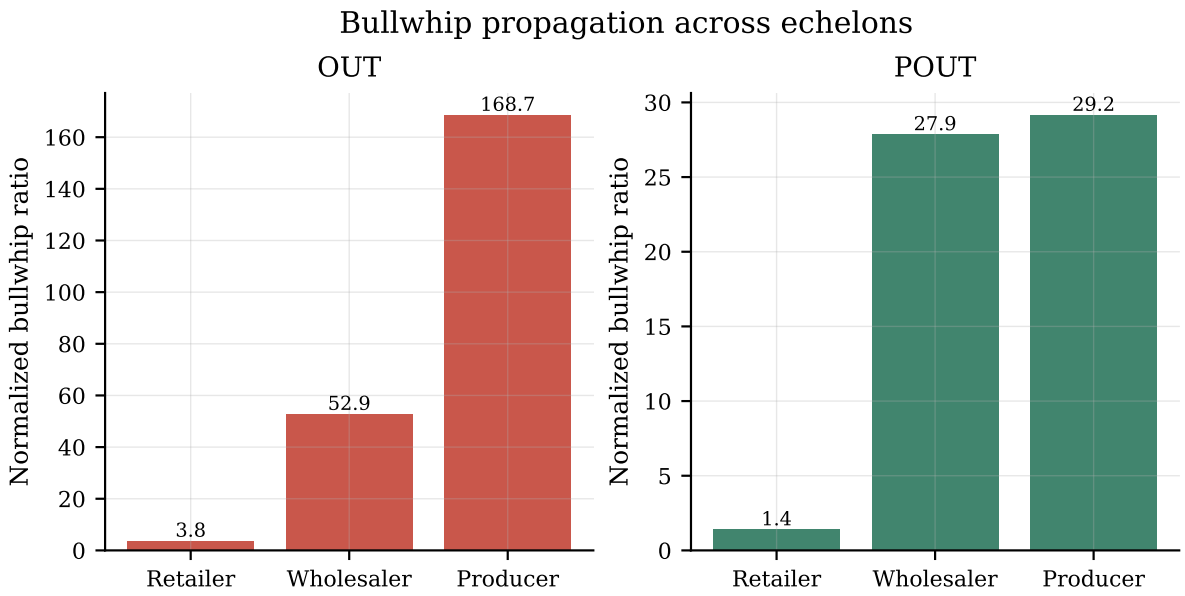


Figure 7: Bullwhip propagation across echelons under each policy at the balanced operating point. The cascade shape is policy-invariant; the level is policy-determined.

Mass balance. In steady state, total production must equal total throughput (demand served), up to warm-up transients. Across configurations the ratio of cumulative production to cumulative demand is 1.0000–1.0008 (for example, 1.0008 at the balanced point under OUT and 1.0000 under POUT), confirming that no material is created or destroyed by the constraint mechanic and that the new-build portion correctly compensates for any recovered shortfall in the mean.

Achieved recovered share. The mass-balance condition Eq. (10) predicts an achieved recovered share of production equal to $\phi\rho$. The simulation reproduces this precisely: at $(\rho, \phi) = (0.3, 0.8)$ the measured share is 0.236–0.240 against a target of 0.240; at $(0.5, 0.6)$ it is 0.300 against a target of 0.300. The recovered share is therefore yield-limited exactly as the analysis predicts, and the constraint is neither slack nor over-binding relative to the mass balance.

Limiting case. Setting $\rho = 0$ yields a recovered share of 0.000 and a constraint frequency of zero, and the producer bullwhip equals the linear-chain value by construction; setting $\phi = 1.0$ drives the constraint frequency to near zero (0.007–0.008 in the full design) and $AR_P \to 1$. Both limits recover the linear chain, as required, and provide a continuous bridge between the closed-loop and linear systems.

Sampling precision. The median relative width of the 95% bootstrap confidence interval on producer bullwhip across the full design is 0.061 (i.e. $\pm 3\%$ of the point estimate), and is tighter under POUT (0.047) than under OUT (0.072), as expected since POUT produces a less heavy-tailed order distribution. The damping effects reported above—an AR_P of 0.77–0.85 at the balanced and low-yield points—are several times larger than this sampling width, so they are not artefacts of replication noise. Thirty replications are sufficient for the qualitative conclusions, and the confidence intervals are reported throughout so the reader can judge precision directly.

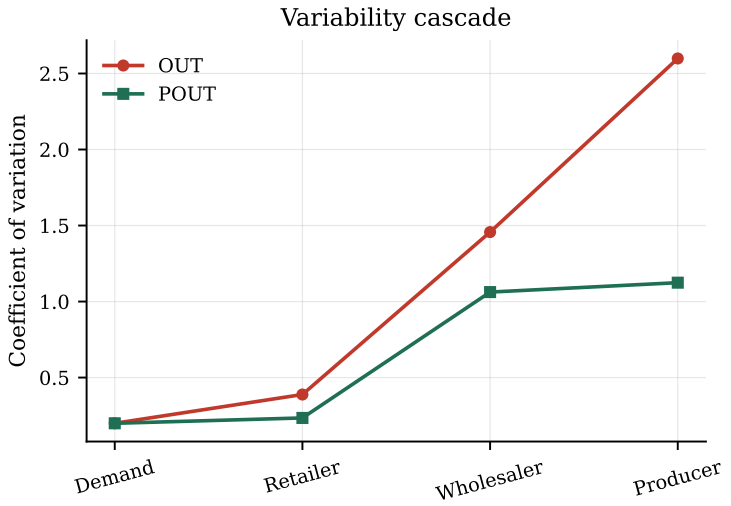


Figure 8: Coefficient-of-variation cascade across echelons, both policies. POUT substantially compresses the upstream growth of variability relative to OUT.

Table 5: Internal-validation diagnostics across the full 576-configuration design (both policies). "Required" is the value implied by the model's construction or by the steady-state analysis of Section 3; "Observed" is the simulated value.

Check	Required	Observed (full design)	Verdict
Production / demand (mass balance)	$= 1$	1.0000–1.0008	pass
Recovered share vs. $\phi\rho$	equal	$r = 0.9995$; mean abs. error 0.0016, max 0.025	pass
Limiting case $\rho = 0$	constraint $= 0$, $AR_P = 1$	constraint $= 0.000$, $AR_P = 1.000$ (by construction)	pass
Limiting case $\phi = 1$	constraint $\to 0$, $AR_P \to 1$	constraint 0.007–0.008 (mean), max 0.015	pass
Sampling precision (rel. 95% CI on BWE_P)	small vs. effect	median 0.061 (OUT 0.072, POUT 0.047); $p_{90} \approx 0.105$	pass

Consolidated checks. Table 5 collects the internal-validation diagnostics in one place. Each row states the quantity checked, the value the model's construction requires, the value the simulation produces across the full design, and the verdict. The model passes every check: it conserves mass, reproduces the analytically predicted recovered share $\phi\rho$ to within 0.16% on average, collapses to the linear chain in both limiting cases ($\rho \to 0$ and $\phi \to 1$), and estimates producer bullwhip with a median relative confidence width of 6.1%—comfortably smaller than the damping effects it is used to detect.

5.8 Factor hierarchy across the full design

The controlled comparison and the mechanism analysis above fix particular operating points. To show that the qualitative ordering of factor importance is not an artefact of those points, Table 6 reports the marginal mean of producer bullwhip at each factor level, averaged over all other factors across the entire design, with 95% confidence intervals. The *range* column—the spread of the marginal mean across a factor's levels—is a simple, design-wide measure of that factor's influence on producer bullwhip.

The hierarchy is unambiguous and consistent across both policies. Under OUT the smoothing pa-

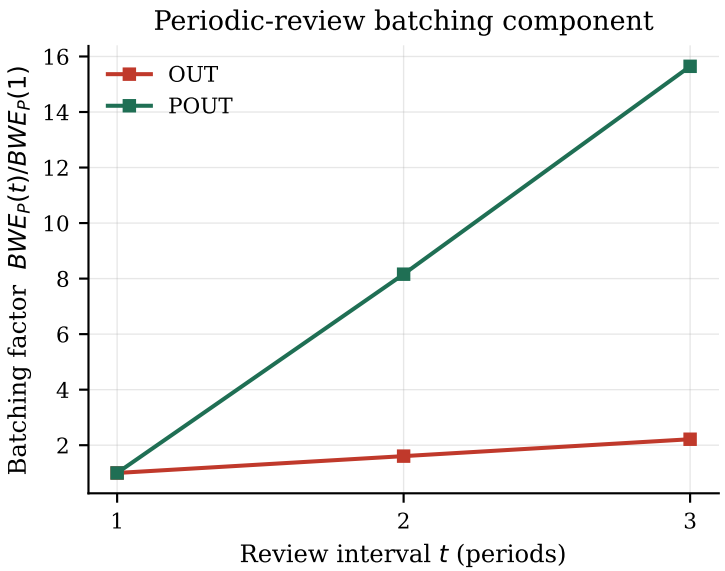


Figure 9: Periodic-review batching factor $\beta_P(t) = BWE_P(t)/BWE_P(1)$ in a near-linear configuration. The review interval is a dominant amplifier of producer bullwhip under both policies.

Table 6: Marginal mean producer bullwhip BWE_P at each factor level, averaged over all other factors across the full design, with 95% confidence intervals. The *range* (max−min across levels) measures design-wide factor influence.

		Marginal mean BWE_P (±95% CI) by level				Range
Policy	Factor					
OUT	α	60.4±5.8	166.3±9.6	262.9±13.6	375.8±23.1	315.4
OUT	t	127.4±11.0	214.9±17.1	306.7±23.6		179.3
OUT	ϕ	190.2±18.7	211.2±20.1	247.6±22.5		57.4
OUT	ρ	236.7±25.7	224.0±24.5	210.1±23.2	194.7±21.9	42.0
OUT	τ	216.1±24.0	216.3±24.0	216.5±24.0	216.5±24.0	0.4
POUT	t	5.7±0.7	33.4±1.0	60.7±1.0		55.0
POUT	α	27.5±3.6	29.9±3.6	33.9±3.7	41.9±3.9	14.3
POUT	ϕ	31.2±3.2	32.4±3.2	36.2±3.5		5.1
POUT	ρ	34.6±3.9	33.7±3.8	32.8±3.8	32.0±3.7	2.7
POUT	τ	33.3±3.8	33.3±3.8	33.3±3.8	33.2±3.8	0.1

rameter α has by far the largest design-wide influence (range 315), followed by the review interval t (range 179); the refurbishment factors ϕ and ρ are intermediate (ranges 57 and 42); and the refurbishment delay τ is negligible (range 0.4, i.e. smaller than the confidence interval on any single level). Under POUT the controller compresses the absolute levels by an order of magnitude and the review interval becomes the leading factor (range 55 versus 14 for α), but τ remains negligible (range 0.1) and ρ near-negligible (range 2.7). The near-zero influence of τ at every level is the design-wide confirmation of the earlier point-wise finding that feasibility and bullwhip are governed by refurbishment *yield*, not *delay*; the monotone increase of BWE_P in α and t matches the classical and control-theoretic expectations of Section 2. Importantly, the two refurbishment factors ρ and ϕ move producer bullwhip in opposite directions—higher mandate ρ *lowers* it (more binding, more truncation) while higher yield ϕ *raises* it back toward the linear level (less binding)—which is exactly the low-pass-filter mechanism expressed in marginal-mean form.

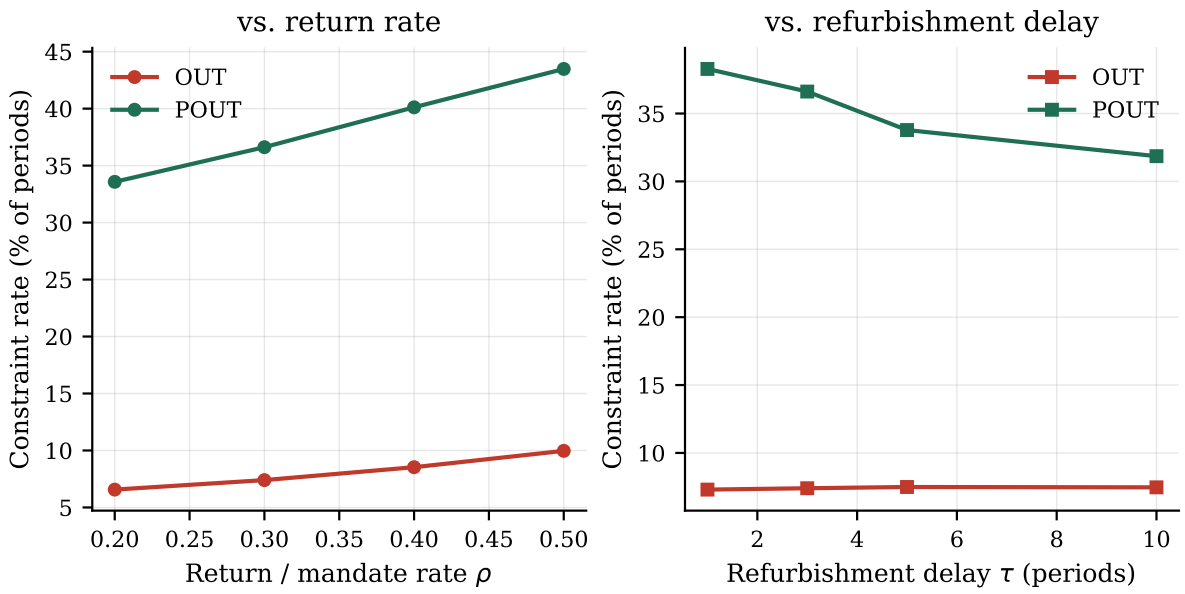


Figure 10: Constraint frequency versus return rate ρ (left) and refurbishment delay τ (right). The mandate rate raises constraint frequency; the delay has only a weak, second-order effect. Yield ϕ (Figure 5) is the dominant lever.

Table 7: Closed-loop damping and constraint frequency by review interval and policy, at $\alpha = 0.4, \rho = 0.3, \tau = 3$. AR_P compares the closed-loop producer bullwhip ($\phi = 0.8$) with the linear baseline ($\phi = 1.0$); lower AR_P means stronger damping.

Policy	Quantity	$t = 1$	$t = 2$	$t = 3$
OUT	AR_P (damping)	0.855	0.869	0.874
OUT	Constraint frequency	0.107	0.074	0.056
POUT	AR_P (damping)	0.719	0.927	0.943
POUT	Constraint frequency	0.633	0.366	0.250

5.9 How the damping interacts with the review interval

Marginal means treat each factor in isolation. The closed-loop damping, however, is most revealing in interaction with the review interval, because the review interval controls how much refurbished buffer the producer accumulates between orders. Table 7 reports, at the balanced point ($\alpha = 0.4, \rho = 0.3, \tau = 3$), the producer amplification ratio AR_P (closed loop at $\phi = 0.8$ relative to the linear $\phi = 1.0$ baseline) and the corresponding constraint frequency, broken out by review interval and policy.

Two patterns stand out. First, under OUT the damping is almost independent of the review interval (AR_P between 0.855 and 0.874): the order-up-to rule issues large, infrequent catch-up pulses whose upper tail is truncated by the constraint to about the same relative degree regardless of how those pulses are batched. Second, under POUT the damping is strongly interval-dependent and is by far deepest at continuous review ($AR_P = 0.719$ at $t = 1$, against 0.943 at $t = 3$). The reason is visible in the constraint-frequency row: at $t = 1$ the proportional controller orders smoothly and frequently, accumulates the thinnest refurbished buffer, and therefore binds in 63% of periods—so the low-pass filter acts on a large fraction of orders and removes substantial variance. As the review interval lengthens, each order draws on buffer accumulated over several periods, the constraint binds less often (25% at $t = 3$), and the

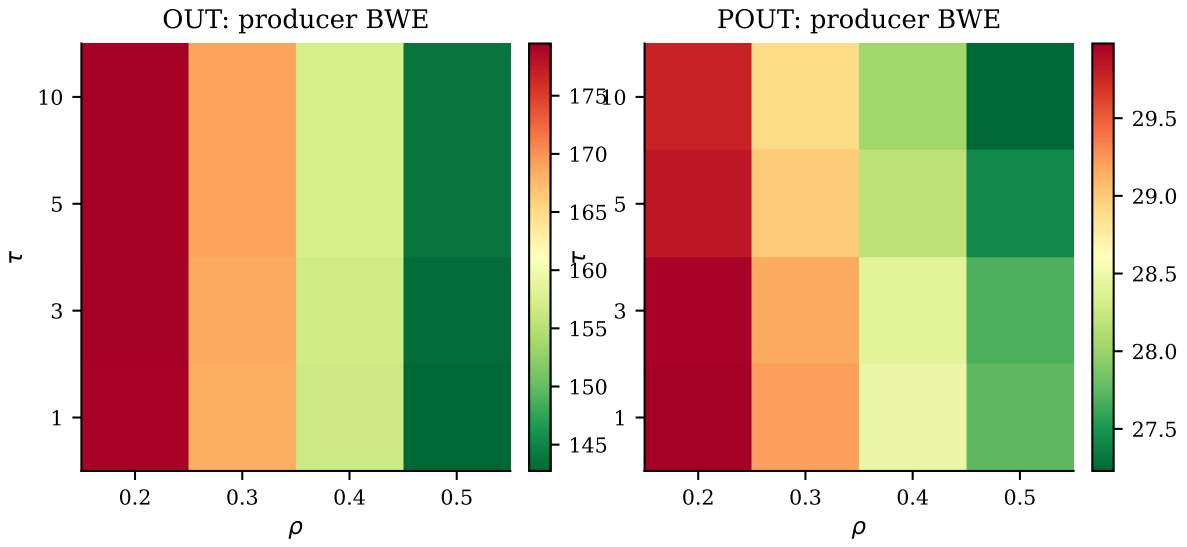


Figure 11: Producer bullwhip over the (ρ, τ) plane for each policy. The low-bullwhip region is substantially larger under POUT.

Table 8: Consolidated results at the balanced operating point ($\alpha = 0.4, t = 2, \rho = 0.3, \tau = 3$). BWE_P is shown with its 95% bootstrap confidence interval; AR_P is the producer amplification ratio relative to the $\phi = 1$ linear-equivalent baseline within each policy. Service level is the fill rate; recovered share is the realised refurbished fraction of production.

Policy	Setting	BWE_R	BWE_W	BWE_P (95% CI)	AR_P	Constr.	Rec. share
OUT	$\phi = 1.0$ (linear-eq.)	3.77	52.85	194.1 [186.7, 202.2]	1.000	0.006	0.295
OUT	$\phi = 0.8$	3.77	52.87	168.7 [162.4, 175.4]	0.869	0.074	0.239
OUT	$\phi = 0.6$	3.77	52.87	154.0 [148.9, 159.7]	0.794	0.131	0.180
POUT	$\phi = 1.0$ (linear-eq.)	1.41	27.87	31.5 [30.7, 32.2]	1.000	0.002	0.300
POUT	$\phi = 0.8$	1.41	27.87	29.2 [28.5, 29.8]	0.927	0.366	0.240
POUT	$\phi = 0.6$	1.41	27.87	28.8 [28.1, 29.4]	0.914	0.484	0.180

damping weakens toward unity. The interaction thus reinforces the central mechanism: the closed-loop damping is not a fixed property of the system but scales with how often the refurbishment constraint actually binds, which the ordering policy and review interval jointly determine.

This also sharpens the policy trade-off of Section 5.6. The combination that minimises absolute producer bullwhip—POUT at continuous review ($BWE_P \approx 2.8$)—is precisely the combination that binds the constraint most often (63%). A manager choosing that operating point for its variance benefits would simultaneously be choosing the regime in which refurbished-component shortfalls are most frequent, and would need to either raise refurbishment yield ϕ or accept the resulting unmet recovered-content target. The variance-optimal and constraint-safe regions of the design do not coincide.

5.10 Summary of results

Table 8 consolidates the headline figures at the balanced operating point in a single reference, spanning both ordering policies, all three echelons, and three refurbishment-yield settings (including the $\phi = 1$ linear-equivalent baseline), each with the 95% confidence interval on producer bullwhip. It is the numerical companion to Figures 1–5 and the source of the amplification ratios quoted throughout.

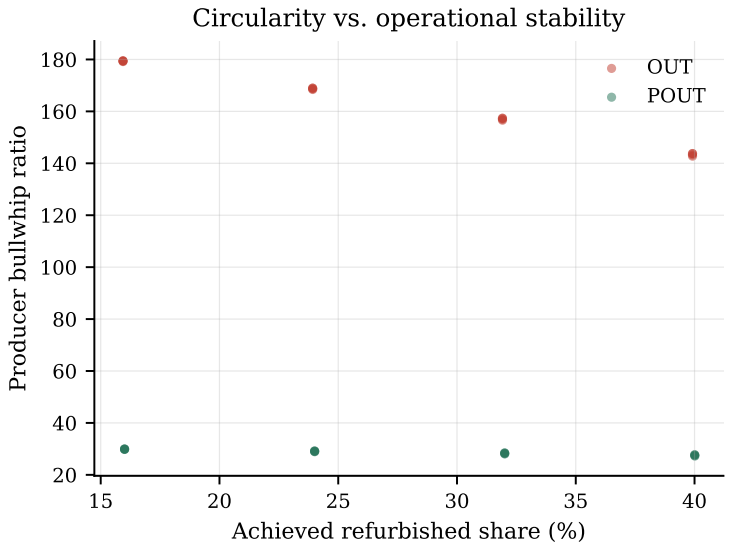


Figure 12: Circularity versus operational stability: producer bullwhip against the achieved recovered share across the design. POUT attains low bullwhip but at higher constraint exposure.

The table makes three things directly checkable. The retailer and wholesaler columns are invariant to the policy's refurbishment setting (the constraint never reaches them); the producer column falls monotonically as ϕ decreases (the deepening low-pass damping); and the recovered-share column tracks $\phi\rho$ exactly ($0.8 \times 0.3 = 0.24$, $0.6 \times 0.3 = 0.18$), while service level remains at 1.000 throughout because the new-build portion always compensates the mean. The constraint frequency rises as ϕ falls and is far higher under POUT, the quantitative form of the policy trade-off.

The results can be summarised in five statements, each supported above. (i) Relative to an identical linear chain, the mandatory refurbishment constraint leaves downstream echelons unchanged and dampens producer bullwhip ($AR_P < 1$). (ii) The damping is a low-pass filtering effect: the constraint truncates the upper tail of the producer order distribution, and deepens as the constraint binds more often. (iii) Absolute bullwhip levels are set primarily by the ordering policy (an order-of-magnitude OUT-versus-POUT difference) and secondarily by the periodic-review interval, the single largest amplifier. (iv) Feasibility is governed by refurbishment yield through $\phi\,\gamma_{ret} \geq 1$, with the achieved recovered share equal to $\phi\rho$; the delay τ is second-order. (v) A policy-dependent trade-off links the variance-damping proportional policy to a higher constraint frequency.

6 Sensitivity Analysis

The headline results are reported at a balanced operating point and over a fixed factorial grid. This section tests their robustness to three modelling choices that are not varied in the main design: the proportional controller horizon T_i, the demand variability CV_D, and the smoothing parameter α examined jointly with the linear-versus-closed-loop contrast. Throughout, we vary one quantity while holding the others at the balanced operating point, and we report producer-level outcomes since the closed-loop effect is localised there.

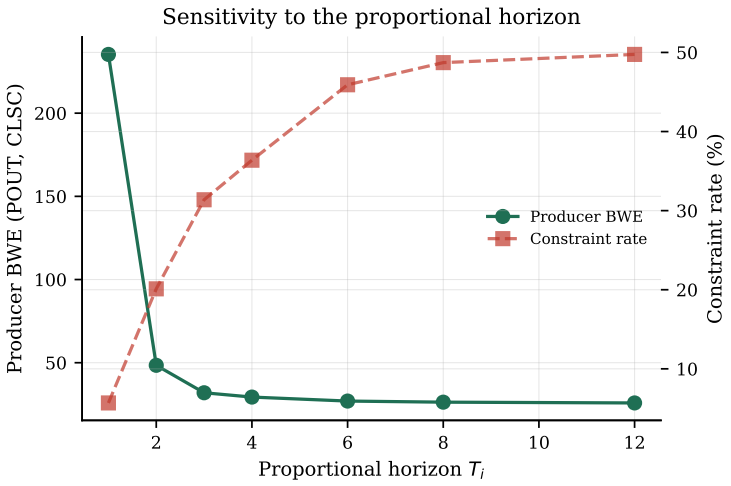


Figure 13: Sensitivity to the proportional horizon T_i (POUT, closed-loop, balanced point): producer bullwhip (solid) and constraint frequency (dashed). Most variance reduction is achieved by $T_i \approx 4$; the constraint frequency rises with T_i.

6.1 Proportional controller horizon

The POUT rule has one free tuning parameter, the horizon T_i, which we fix at 4 in the main study. Figure 13 traces producer bullwhip and the constraint frequency as T_i ranges from 1 (where POUT coincides with OUT) to 12. Producer bullwhip falls monotonically and steeply as T_i increases from 1, then flattens: most of the variance reduction is achieved by $T_i \approx 4$, beyond which further smoothing yields diminishing returns. The constraint frequency rises with T_i, confirming the policy-dependent trade-off of Section 5.6 as a continuous effect rather than a two-point artefact: smoother ordering (higher T_i) lowers bullwhip but accumulates less recovered buffer and binds more often. The choice $T_i = 4$ thus sits in a sensible region—near the knee of the bullwhip curve, before the constraint frequency grows unacceptably—and the qualitative results are insensitive to moderate variation around it.

6.2 Demand variability

Figure 14 varies the demand coefficient of variation CV_D from 0.10 to 0.40 and shows producer bullwhip for the linear and closed-loop systems under each policy. The normalised bullwhip ratio is, as expected from its scale-invariant construction, only weakly dependent on CV_D: the curves are nearly flat, confirming that Eq. (11) measures relative amplification rather than absolute dispersion. Importantly, the closed-loop curve lies at or below the linear curve across the whole range under both policies, so the localised damping result is not an artefact of the particular $CV_D = 0.20$ used in the main design; it holds for both calmer and more volatile demand.

6.3 Smoothing parameter and the closed-loop contrast

Figure 15 shows producer bullwhip as a function of the smoothing parameter α, overlaying the linear and closed-loop systems for each policy. Two features are robust across α. First, producer bullwhip increases with α under both policies and both systems, reflecting that a more responsive forecast chases noise and inflates the order signal; the effect is far stronger under OUT than under POUT, whose proportional feed-

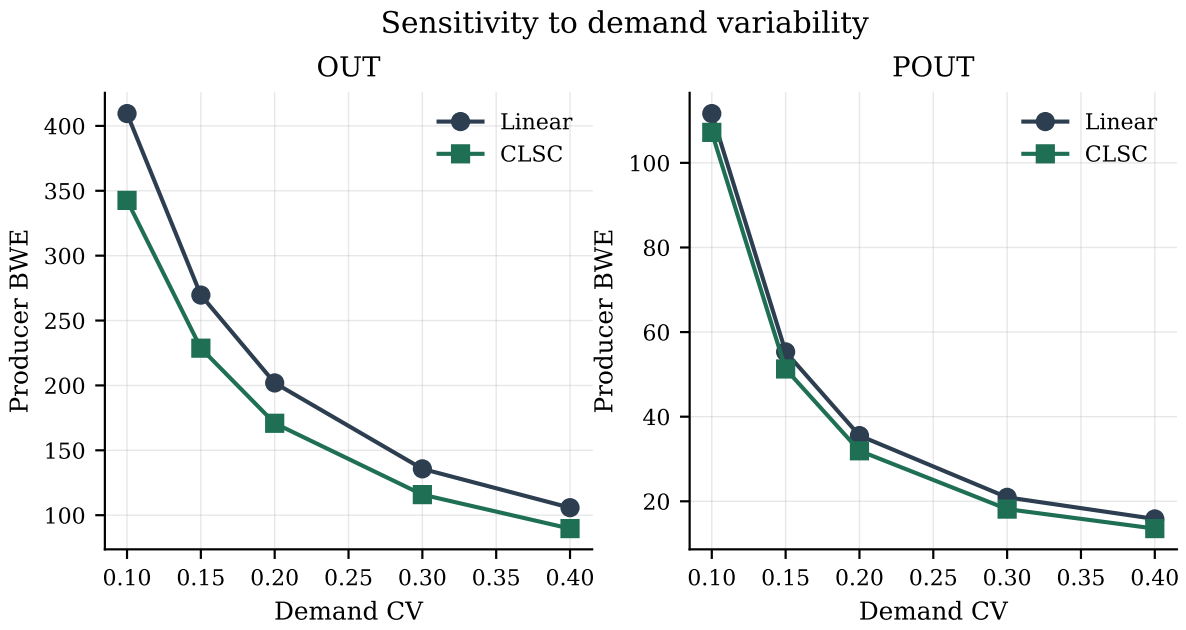


Figure 14: Sensitivity to demand variability CV_D: producer bullwhip, linear versus closed-loop, each policy. The normalised ratio is nearly invariant to CV_D, and the closed-loop damping holds across the range.

back attenuates the additional variance. Second, the closed-loop curve tracks at or below the linear curve for every value of α, so the damping result is not specific to the $\alpha = 0.4$ used elsewhere. The vertical gap between the linear and closed-loop curves—the absolute damping—widens modestly at higher α under OUT, because a more aggressive forecast produces larger catch-up releases for the constraint to clip.

6.4 Refurbishment delay revisited

Table 9 reports the producer bullwhip and constraint frequency as the refurbishment delay τ varies, complementing Figure 10. Under OUT the producer bullwhip and constraint frequency are essentially flat in τ; under POUT the constraint frequency declines slightly as τ grows, because a longer pipeline smooths the timing of recovered arrivals, but the producer bullwhip is again nearly unchanged. This confirms quantitatively that τ is a second-order lever relative to the yield ϕ: its main effect is on the timing, not the rate, of recovered supply, and the mean-rate-driven mass-balance condition Eq. (10) is what governs feasibility.

6.5 Summary of sensitivity

The five qualitative findings of Section 5.10 survive variation in the proportional horizon, the demand variability, the smoothing parameter and the refurbishment delay. The absolute levels move with T_i and α, as expected, but the localised damping, the order-of-magnitude dominance of ordering policy, and the yield-governed feasibility all persist. This supports the scope statement of Section 7.7: the qualitative conclusions follow from the structure of the model rather than from the particular parameter values chosen.

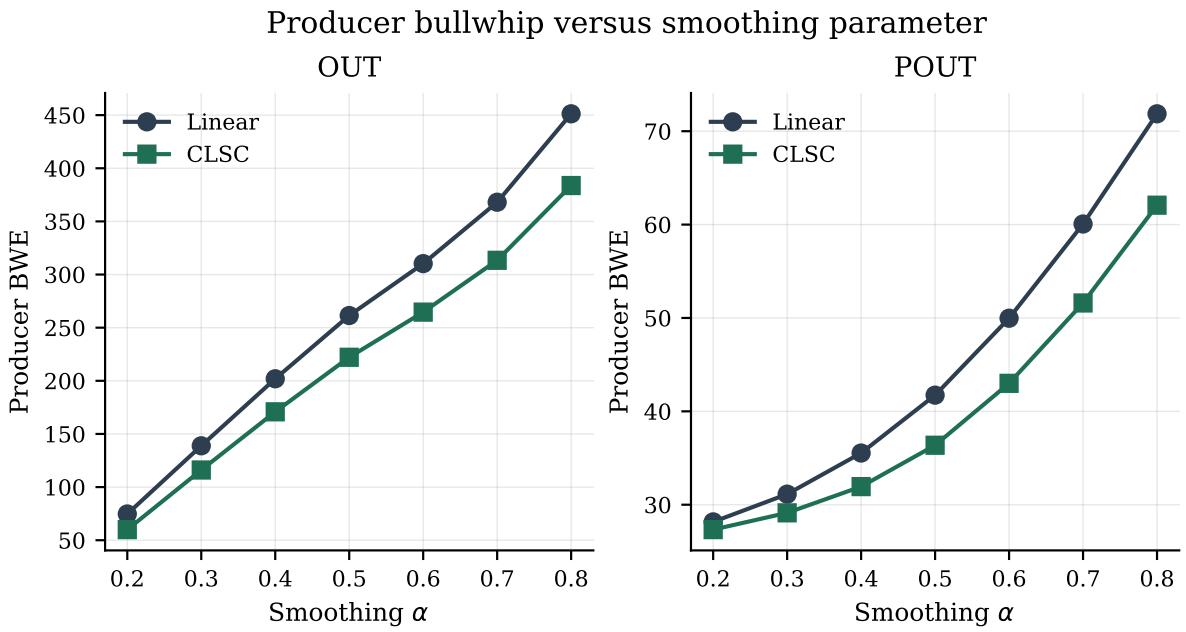


Figure 15: Producer bullwhip versus smoothing parameter α, linear versus closed-loop, each policy. The closed-loop curve lies at or below the linear curve throughout, confirming the damping result is robust in α.

Table 9: Sensitivity to the refurbishment delay τ (balanced point except τ). Producer bullwhip and constraint frequency are nearly flat in τ under OUT; under POUT the constraint frequency declines slightly with τ.

Policy	τ	BWE_P	Constraint rate
OUT	1	170.62	0.071
	3	170.83	0.072
	5	171.12	0.072
	10	171.17	0.072
POUT	1	32.10	0.339
	3	31.94	0.314
	5	31.78	0.288
	10	31.84	0.287

7 Discussion

7.1 Does the constraint amplify the bullwhip effect?

The controlled comparison answers the paper's motivating question directly: relative to an otherwise identical linear supply chain, a mandatory refurbishment constraint does not amplify the bullwhip effect at the operating points studied. Downstream echelons are unaffected, and the producer is mildly damped, with the damping strengthening as the constraint binds. For circular-economy policy this is a reassuring conclusion. A frequent concern about recovered-content mandates is that they might destabilise production planning by introducing a second, uncertain supply source subject to its own delays and yields. Our results suggest that, in the binding-constraint regime, the mandate does not add demand-variability amplification on top of what the forward chain already generates; the dominant sources of volatility—the ordering policy and the review interval—are properties of the forward system that the mandate does not

worsen.

This should not be read as a claim that recovered-content mandates are operationally costless. The cost surfaces elsewhere, in the constraint frequency: when yield is low the producer is frequently unable to meet its desired release from recovered stock, which in a richer model would translate into unmet output, expediting, or quota non-compliance. The bullwhip ratio improves precisely because those large releases are suppressed. The right reading is therefore that the mandate converts a variability problem into a service/compliance problem at the producer, rather than propagating variability through the chain.

7.2 The constraint as a filter, not a destabiliser

The damping mechanism is worth dwelling on because it inverts a natural intuition. A binding capacity-like constraint is often expected to destabilise, by forcing catch-up behaviour once the constraint relaxes. In our model the catch-up does occur—backlog accumulates when the constraint binds—but the measured *order* variance nonetheless falls, because the constraint removes the largest upward production releases from the realised order stream (Figures 3–4). The bullwhip ratio is a property of the order distribution, and clipping its upper tail reduces its variance. This is a genuinely different phenomenon from the stabilisation achieved by a proportional controller: the controller *chooses* to issue smaller orders to reduce variance, whereas the constraint *prevents* the largest orders from being executed. The two can co-exist, and indeed POUT plus a binding constraint yields both the lowest absolute bullwhip and the highest constraint frequency in our design. The same nonlinearity has been observed in forward chains: Ponte et al. (2017) show that capacity constraints reshape—and can reduce—bullwhip in open supply chains, and Dominguez et al. (2019) report the analogous dampening in closed loops. Our contribution is to show that a constraint acting on *only the recovered-content portion* of production, rather than on total throughput, produces the same qualitative filtering, and to localise it precisely at the producer while leaving the downstream echelons untouched. The describing-function analyses of nonlinear production-inventory systems by Spiegler et al. (2016) provide the methodological frame within which such truncation effects are formally studied.

This mechanism also supplies the explicit explanation that the review-level literature has lacked. Systematic reviews (Braz et al., 2018) report that closing a supply chain can reduce the bullwhip effect and attribute the reduction loosely to higher return rates. Our analysis localises the reduction to the constrained echelon and identifies its precise cause as the truncation of the order distribution's upper tail, with a clean monotone dependence on the binding frequency. It also clarifies why the literature reports conflicting signs: whether the closed loop amplifies or dampens depends on whether the recovered flow is modelled as a binding production constraint (our regime, which dampens at the producer) or as a planning-visible buffer that is anticipated in the ordering decision (which would change the order signal itself). The two regimes are different couplings of the same physical return flow, and they need not share a sign—consistent with the visibility-dependence emphasised by Ponte et al. (2020).

7.3 Ordering policy as the primary operational lever

Across the entire design, the choice between the order-up-to and proportional order-up-to rules changes bullwhip outcomes by roughly an order of magnitude—far more than the closed-loop coupling and far

more than any single physical parameter such as the return rate, the delay, or the yield. For a firm operating under a refurbishment mandate, the practical implication is that the most effective and least expensive bullwhip intervention is the structure of the ordering controller, not the fine-tuning of the refurbishment parameters, which are often outside the firm's short-run control (yield depends on the quality of the return stream; delay on the refurbishment technology). The proportional controller is, moreover, the control structure established as appropriate for closed-loop chains by Cannella et al. (2021), so its adoption does not require abandoning the closed-loop setting; it is complementary to it. The caveat, developed in Section 5.6, is that the proportional policy raises the constraint frequency, so its adoption should be paired with recovered-stock buffering or yield improvement.

7.4 Yield-governed feasibility and its managerial reading

The empirical collapse of the constraint at $\phi = 1$ and its severity at $\phi = 0.6$ confirm that feasibility is a yield phenomenon governed by $\phi\,\gamma_{ret} \geq 1$. The managerial reading is sharp: investment in refurbishment *yield*—the reliability of the recovery process, the quality of inbound returns, the effectiveness of inspection—relaxes binding constraints far more effectively than investment in refurbishment *speed*. A faster refurbishment process (lower τ) changes the timing of recovered arrivals but not their long-run rate, and so does little to ease a binding mandate; a higher-yield process changes the rate directly. This places our feasibility result alongside the yield rate paradox of Hosoda et al. (2021): where they show that recovery yield governs closed-loop dynamics in non-obvious ways, we show that it also governs the hard feasibility boundary of a mandated quota. Equally, raising the mandated recovered content ρ without a commensurate improvement in yield is what makes the constraint bind, and pushes the system toward the high-constraint region of the design. Policy that mandates recovered content should therefore be paired with measures that raise achievable yield—design-for-disassembly, return-quality incentives, inspection technology—if the intent is to avoid chronic under-fulfilment of the quota.

7.5 Positioning against specific prior findings

It is useful to place each result against the specific prior findings it bears on. Zhou and Disney (2006) establish that returns can either amplify or dampen order and inventory variance depending on the return rate and the remanufacturing lead time. Our results refine this for the binding-constraint regime: within that regime the return rate (mandate rate ρ) acts only in the damping direction at the producer—higher ρ produces more clipping of the order tail and slightly lower producer bullwhip (the marginal effect of ρ on BWE_P is weakly negative in our full design, Appendix B)—and the delay is second-order, consistent with their identification of lead time and return rate as the two governing quantities while sharpening which direction each pushes under a binding mandate. Adenso-Díaz et al. (2012) single out the return rate as a mitigating factor in reverse chains; our tail-truncation mechanism is a concrete channel through which a higher mandated return rate mitigates producer variability. Dominguez et al. (2020) show that remanufacturing lead-time *variability* is the damaging quantity; our model uses a deterministic delay and accordingly finds the (mean) delay to be second-order, which is consistent with their result that it is the variability, not the mean, of the delay that matters—an explicit avenue for extending our model. Ponte et al. (2020) demonstrate that the sign and magnitude of return-rate and lead-time effects depend on information transparency; our binding-constraint regime corresponds to the low-visibility case in which

recovered supply is not anticipated in the ordering decision, and the contrasting buffer regime we describe in Section 7 corresponds to their high-visibility case, so our two regimes can be read as endpoints of their visibility spectrum. Finally, the systematic review of Braz et al. (2018) reports that closing a chain can reduce the bullwhip effect and attributes this loosely to higher return rates; our analysis supplies the missing mechanism (upper-tail truncation at the constrained echelon) and the missing locus (the producer alone).

7.6 Managerial and policy implications

The results translate into a small number of concrete recommendations, ordered by the size of their effect in our design.

For operations managers. The single largest bullwhip lever is the ordering controller. Replacing a classical order-up-to rule with a proportional order-up-to rule reduces producer bullwhip by roughly an order of magnitude in our design—larger than the effect of any physical parameter and larger than the closed-loop coupling itself. The second lever is the review cadence: because the review interval is the dominant batching amplifier, reviewing more frequently (or, where review frequency is fixed by system constraints, reporting a batching-corrected bullwhip metric to avoid mis-attributing batching to feedback) materially reduces measured producer variability. The third lever, specific to the closed loop, is recovered-stock buffering: a manager who adopts the proportional controller for its variance benefit should anticipate a higher constraint frequency and buffer recovered stock accordingly, since the smoother order stream accumulates less recovered inventory between binding episodes.

For circular-economy policymakers. The analysis is reassuring on one front and cautionary on another. Reassuringly, a mandatory recovered-content quota does not, in the binding-constraint regime, amplify demand-variability through the chain; the concern that recovered-content mandates destabilise upstream production planning is not borne out here, and the mandate's effect on order variability is, if anything, stabilising at the producer. Cautiously, however, the operational burden of the mandate appears not as variability but as constraint frequency: when refurbishment yield is low, the producer is frequently unable to meet the mandated recovered content, which in practice would manifest as quota non-compliance, expediting, or unmet output. The policy lever that addresses this is yield, not speed: because feasibility is governed by $\phi\,\gamma_{ret} \geq 1$, measures that raise achievable yield—design-for-disassembly requirements, return-quality and collection incentives, inspection and grading technology—relax binding mandates far more effectively than measures that merely accelerate refurbishment. A mandate that raises the required recovered content ρ without a commensurate yield improvement predictably pushes producers into the chronic-binding region. This gives a quantitative rationale for pairing recovered-content targets with yield-enabling design and collection requirements, rather than imposing content targets in isolation.

7.7 Scope and limitations

The findings hold within the modelling assumptions stated in Section 3, and several of these bound the generality of the conclusions. The demand process is stationary, normally distributed and serially inde-

pendent; autocorrelated or trended demand would change the quantitative bullwhip levels and, as Gaalman et al. (2022) show, can even change the direction in which bullwhip moves with the lead time. The chain is serial and four-echelon; divergent or networked structures with multiple collectors are known to behave differently. The constraint mechanic is a specific binding-from-below rule on a recovered-content quota; alternative formulations—a planning-visible recovered buffer anticipated in the order decision, or a soft penalty rather than a hard cap—would alter and possibly reverse the sign of the closed-loop effect, as discussed above. The proportional controller horizon is fixed at $T_i = 4$; the absolute POUT levels depend on this choice, although the qualitative contrast with OUT does not. Finally, the model is calibration-light and expressed in abstract units; absolute bullwhip magnitudes are therefore model artefacts and should not be compared directly with empirically calibrated studies. What is robust is the set of *qualitative* results: the localisation of the closed-loop effect at the producer, its damping sign and yield-deepening behaviour, the order-of-magnitude dominance of ordering policy, the batching dominance of the review interval, and the yield-governed feasibility condition. These follow from the structure of the model rather than from particular parameter values, and the 95% confidence intervals confirm they are not artefacts of sampling variability.

7.8 Implications for research and practice

For research, the controlled-comparison design offers a template for isolating the effect of a single closed-loop coupling without the confounds that arise when results are compared across separately calibrated models; the conflicting signs in the closed-loop bullwhip literature are at least partly attributable to such confounds, and a common model with a toggled coupling is one way to dissolve them. The explicit distributional mechanism—tail truncation—also suggests that future closed-loop bullwhip studies should report order *distributions*, not only variance ratios, since two systems with similar variance can have very different tail behaviour and hence very different service and compliance implications.

For practice, the results point to three levers in priority order. First, the ordering policy: adopting a proportional controller is the largest single bullwhip reduction available, an order of magnitude larger than the closed-loop effect. Second, the review cadence: more frequent review removes the batching amplification that dominates producer variability. Third, refurbishment yield: it governs whether the recovered-content mandate can be met without chronic binding, and dominates refurbishment speed in that role. A firm that sequences its investments in this order will address the largest sources of volatility and infeasibility first.

8 Conclusion

We have compared a closed-loop supply chain with a mandatory refurbishment constraint against an otherwise identical classical linear supply chain, using a single discrete-event Monte-Carlo model in which the linear chain is the $\rho = 0$ special case. The comparison is fully controlled: the two systems share the demand process, ordering policy, lead times and review schedule, and differ only in the mandatory-refurbishment coupling at the producer. Over a full-factorial design of 576 configurations replicated 30 times—34,560 runs across two ordering policies, with bullwhip measured from simulated order streams and reported with 95% confidence intervals—five robust findings emerge.

First, relative to the linear chain the constraint leaves downstream echelons statistically unchanged and *dampens* producer bullwhip; the closed-loop effect is localised entirely at the constrained echelon and does not amplify. Second, the damping is a low-pass filtering effect: the binding constraint truncates the upper tail of the producer order distribution, and the damping deepens monotonically as the constraint binds more often, recovering the linear chain exactly when the constraint never binds. Third, absolute bullwhip levels are set primarily by the ordering policy—an order-of-magnitude difference between the order-up-to and proportional order-up-to rules—and secondarily by the periodic-review interval, which is the single largest amplifier of producer variability. Fourth, system feasibility is governed by refurbishment yield through the dimensionally consistent mass-balance condition $\phi\,\gamma_{ret} \geq 1$, with the achieved recovered share equal to $\phi\rho$ and the delay τ playing only a second-order role. Fifth, a policy-dependent trade-off links the variance-damping proportional policy to a higher constraint frequency, because smoother ordering accumulates less recovered buffer.

Taken together, these results give an explicit mechanism for the recurring but loosely explained observation that closing a supply chain can reduce variability, and they reposition ordering-policy choice—rather than the tuning of refurbishment parameters—as the primary operational lever for stable circular supply chains. They also clarify why the closed-loop bullwhip literature reports conflicting signs: the sign depends on whether the recovered flow is modelled as a binding production constraint or as a planning-visible buffer.

Future work should pursue three directions. The most important is to embed the constraint and buffer couplings in a single parameterised family—for instance by a visibility parameter that interpolates between a fully binding, unanticipated constraint and a fully planning-visible buffer—so that the sign and magnitude of the closed-loop bullwhip effect can be traced continuously between the two regimes and the conflicting literature reconciled within one model. A second direction is to relax the stationary, serially independent demand assumption and to test whether the tail-truncation mechanism survives under autocorrelated and trended demand. A third is to validate the yield-governed feasibility condition and the tail-truncation mechanism against industry return-flow and order data, for which the WEEE and broader EPR reporting streams are a natural empirical setting.

Reproducibility

All results are produced by a self-contained discrete-event simulation. The core model, the full-factorial experiment runner, the controlled linear-versus-CLSC comparison, and the figure-generation scripts are provided as supplementary material, together with the aggregated results files (per-cell means and 95% bootstrap confidence intervals). Each run is seeded, so the entire study is exactly reproducible by re-executing the experiment and figure scripts in sequence.

References

Adenso-Díaz, B., Moreno, P., Gutiérrez, E., Lozano, S., 2012. An analysis of the main factors affecting bullwhip in reverse supply chains. *International Journal of Production Economics* 135(2), 917–928.

Atasu, A., Sarvary, M., Van Wassenhove, L.N., 2008. Remanufacturing as a marketing strategy. *Management Science* 54(10), 1731–1746.

Baldé, C.P., Kuehr, R., Yamamoto, T., et al., 2024. *The Global E-waste Monitor 2024*. International Telecommunication Union (ITU) and United Nations Institute for Training and Research (UNITAR), Geneva/Bonn.

Batista, L., Bourlakis, M., Smart, P., Maull, R., 2018. In search of a circular supply chain archetype. *Production Planning & Control* 29(6), 438–451.

Braz, A.C., De Mello, A.M., de Vasconcelos Gomes, L.A., de Souza Nascimento, P.T., 2018. The bullwhip effect in closed-loop supply chains: A systematic literature review. *Journal of Cleaner Production* 202, 376–389.

Cachon, G.P., Randall, T., Schmidt, G.M., 2007. In search of the bullwhip effect. *Manufacturing & Service Operations Management* 9(4), 457–479.

Cannella, S., Bruccoleri, M., Framinan, J.M., 2016. Closed-loop supply chains: What reverse logistics factors influence performance? *International Journal of Production Economics* 175, 35–49.

Cannella, S., Dominguez, R., Ponte, B., Framinan, J.M., 2018. Capacity restrictions and supply chain performance: Modelling and analysing load-dependent lead times. *International Journal of Production Economics* 204, 264–277.

Cannella, S., Ponte, B., Dominguez, R., Framinan, J.M., 2021. Proportional order-up-to policies for closed-loop supply chains: The dynamic effects of inventory controllers. *International Journal of Production Research* 59(11), 3323–3337.

Chen, Y.F., Disney, S.M., 2007. The myopic order-up-to policy with a proportional feedback controller. *International Journal of Production Research* 45(2), 351–368.

Chen, F., Drezner, Z., Ryan, J.K., Simchi-Levi, D., 2000. Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Science* 46(3), 436–443.

Croson, R., Donohue, K., 2006. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science* 52(3), 323–336.

Dejonckheere, J., Disney, S.M., Lambrecht, M.R., Towill, D.R., 2003. Measuring and avoiding the bullwhip effect: A control theoretic approach. *European Journal of Operational Research* 147(3), 567–590.

Dejonckheere, J., Disney, S.M., Lambrecht, M.R., Towill, D.R., 2004. The impact of information enrichment on the bullwhip effect in supply chains: A control engineering perspective. *European Journal of Operational Research* 153(3), 727–750.

Disney, S.M., Lambrecht, M.R., 2008. On replenishment rules, forecasting, and the bullwhip effect in supply chains. *Foundations and Trends in Technology, Information and Operations Management* 2(1), 1–80.

Disney, S.M., Towill, D.R., 2003. On the bullwhip and inventory variance produced by an ordering policy. *Omega* 31(3), 157–167.

Disney, S.M., Farasyn, I., Lambrecht, M.R., Towill, D.R., Van de Velde, W., 2006. Taming the bullwhip effect whilst watching customer service in a single supply chain echelon. *European Journal of Operational Research* 173(1), 151–172.

Dominguez, R., Ponte, B., Cannella, S., Framinan, J.M., 2019. On the dynamics of closed-loop supply chains with capacity constraints. *Computers & Industrial Engineering* 128, 91–103.

Dominguez, R., Cannella, S., Ponte, B., Framinan, J.M., 2020. On the dynamics of closed-loop supply chains under remanufacturing lead time variability. *Omega* 97, 102106.

Forrester, J.W., 1961. *Industrial Dynamics*. MIT Press, Cambridge, MA.

Fussone, R., Dominguez, R., Cannella, S., Framinan, J.M., 2024. Bullwhip effect in closed-loop supply chains with multiple reverse flows: A simulation study. *Flexible Services and Manufacturing Journal* 36(1), 250–278.

Gaalman, G., Disney, S.M., Wang, X., 2022. When bullwhip increases in the lead time: An eigenvalue analysis of ARMA demand. *International Journal of Production Economics* 250, 108623.

Geissdoerfer, M., Savaget, P., Bocken, N.M.P., Hultink, E.J., 2017. The circular economy — A new sustainability paradigm? *Journal of Cleaner Production* 143, 757–768.

Goltsos, T.E., Ponte, B., Wang, S., Liu, Y., Naim, M.M., Syntetos, A.A., 2019. The boomerang returns? Accounting for the impact of uncertainties on the dynamics of remanufacturing systems. *International Journal of Production Research* 57(23), 7361–7394.

Govindan, K., Hasanagic, M., 2018. A systematic review on drivers, barriers, and practices towards circular economy: A supply chain perspective. *International Journal of Production Research* 56(1–2), 278–311.

Guide, V.D.R., Jr., Van Wassenhove, L.N., 2009. The evolution of closed-loop supply chain research. *Operations Research* 57(1), 10–18.

Hosoda, T., Disney, S.M., 2018. A unified theory of the dynamics of closed-loop supply chains. *European Journal of Operational Research* 269(1), 313–326.

Hosoda, T., Disney, S.M., Gavirneni, S., 2015. The impact of information sharing, random yield, correlation, and lead times in closed loop supply chains. *European Journal of Operational Research* 246(3), 827–836.

Hosoda, T., Disney, S.M., Zhou, L., 2021. The yield rate paradox in closed-loop supply chains. *International Journal of Production Economics* 239, 108187.

Kirchherr, J., Reike, D., Hekkert, M., 2017. Conceptualizing the circular economy: An analysis of 114 definitions. *Resources, Conservation and Recycling* 127, 221–232.

Lee, H.L., Padmanabhan, V., Whang, S., 1997. Information distortion in a supply chain: The bullwhip effect. *Management Science* 43(4), 546–558.

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

Lin, J., Naim, M.M., Purvis, L., Gosling, J., 2017. The extension and exploitation of the inventory and order based production control system archetype from 1982 to 2015. *International Journal of Production Economics* 194, 135–152.

Lin, J., Zhou, L., Spiegler, V.L.M., Naim, M.M., Syntetos, A., 2022. Push or pull? The impact of ordering policy choice on the dynamics of a hybrid closed-loop supply chain. *European Journal of Operational Research* 300(1), 282–295.

Morseletto, P., 2020. Targets for a circular economy. *Resources, Conservation and Recycling* 153,

104553.

Ponte, B., Wang, X., de la Fuente, D., Disney, S.M., 2017. Exploring nonlinear supply chains: The dynamics of capacity constraints. *International Journal of Production Research* 55(14), 4053–4067.

Ponte, B., Naim, M.M., Syntetos, A.A., 2019. The value of regulating returns for enhancing the dynamic behaviour of hybrid manufacturing-remanufacturing systems. *European Journal of Operational Research* 278(2), 629–645.

Ponte, B., Framinan, J.M., Cannella, S., Dominguez, R., 2020. Quantifying the Bullwhip Effect in closed-loop supply chains: The interplay of information transparencies, return rates, and lead times. *International Journal of Production Economics* 230, 107798.

Ponte, B., Dominguez, R., Cannella, S., Framinan, J.M., 2022. The implications of batching in the bullwhip effect and customer service of closed-loop supply chains. *International Journal of Production Economics* 244, 108379.

Potting, J., Hekkert, M., Worrell, E., Hanemaaijer, A., 2017. *Circular Economy: Measuring Innovation in the Product Chain*. PBL Netherlands Environmental Assessment Agency, The Hague.

Reike, D., Vermeulen, W.J.V., Witjes, S., 2018. The circular economy: New or refurbished as CE 3.0? — Exploring controversies in the conceptualization of the circular economy through a focus on history and resource value retention options. *Resources, Conservation and Recycling* 135, 246–264.

Souza, G.C., 2013. Closed-loop supply chains: A critical review, and future research. *Decision Sciences* 44(1), 7–38.

Spiegler, V.L.M., Potter, A.T., Naim, M.M., Towill, D.R., 2016. The value of nonlinear control theory in investigating the underlying dynamics and resilience of a grocery supply chain. *International Journal of Production Research* 54(1), 265–286.

Sterman, J.D., 1989. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision-making experiment. *Management Science* 35(3), 321–339.

Tang, O., Naim, M.M., 2004. The impact of information transparency on the dynamic behaviour of a hybrid manufacturing/remanufacturing system. *International Journal of Production Research* 42(19), 4135–4152.

Turrisi, M., Bruccoleri, M., Cannella, S., 2013. Impact of reverse logistics on supply chain performance. *International Journal of Physical Distribution & Logistics Management* 43(7), 564–585.

van der Laan, E., Salomon, M., Dekker, R., Van Wassenhove, L., 1999. Inventory control in hybrid systems with remanufacturing. *Management Science* 45(5), 733–747.

Wang, X., Disney, S.M., 2016. The bullwhip effect: Progress, trends and directions. *European Journal of Operational Research* 250(3), 691–701.

Zhou, L., Disney, S.M., 2006. Bullwhip and inventory variance in a closed loop supply chain. *OR Spectrum* 28(1), 127–149.

Zhou, L., Naim, M.M., Disney, S.M., 2017. The impact of product returns and remanufacturing uncertainties on the dynamic performance of a multi-echelon closed-loop supply chain. *International Journal of Production Economics* 183(B), 487–502.

A Simulation algorithm

Algorithm 10 states the per-period simulation logic in pseudocode, corresponding to the event sequence of Section 3.8. The recovered pipeline is a first-in-first-out queue of length τ; returns enter at its tail and emerge from its head at yield ϕ. Each forward echelon holds an analogous lead-time pipeline. The linear chain is obtained by setting $\rho = 0$, which makes the recovered requirement and the constraint vanish identically.

Table 10: Per-period simulation logic (one replication).

for each period $t = 1, \ldots, T$:
  draw $D_t \sim \max\{0, \mathcal{N}(\mu_D, \sigma_D^2)\}$
  // Retailer
  update $\hat{\mu}_R, \hat{\sigma}_R$ from D_t; receive pipeline arrival; serve D_t (backlog shortfall)
  if $t \bmod t_R = 0$: place $O_{R,t}$ via OUT (Eq. 6) or POUT (Eq. 7)
  // Wholesaler
  update $\hat{\mu}_W, \hat{\sigma}_W$ from $O_{R,t}$; receive arrival; ship to R; **if** review: place $O_{W,t}$
  // Producer
  update $\hat{\mu}_P, \hat{\sigma}_P$ from $O_{W,t}$; receive WIP arrival
  advance recovered pipeline: $I^{\text{ref}} \mathrel{+}= \phi \cdot$ (emerging returns); enqueue $\rho \cdot$ (sales)
  ship finished goods to W from on-hand inventory
  if $t \bmod t_P = 0$:
    compute desired release $O_{P,t}^{\text{rel}}$
    $r^{\text{need}} \leftarrow \rho\, O_{P,t}^{\text{rel}};\ n^{\text{need}} \leftarrow (1-\rho)\, O_{P,t}^{\text{rel}}$
    $r^{\text{used}} \leftarrow \min(I^{\text{ref}}, r^{\text{need}});\ Q^{\text{act}} \leftarrow r^{\text{used}} + n^{\text{need}};\ I^{\text{ref}} \mathrel{-}= r^{\text{used}}$
    enqueue Q^{act} into WIP pipeline (arrives after L_P)
  record orders, inventories, I^{ref}, constraint indicator, production split
discard warm-up periods; compute BWE_i (Eq. 11) over the remainder

B Per-factor marginal results

Table 11 reports the marginal means of producer bullwhip and constraint frequency for each factor level, averaged over all other factors and over 30 replications, for both policies. These marginals support the statements in the main text: producer bullwhip increases strongly with α and with the review interval t, decreases weakly with the mandate rate ρ (more clipping), is essentially flat in the delay τ, and increases with the yield ϕ (a higher yield means fewer binding episodes and therefore less tail truncation). The constraint frequency decreases with α, with t, and sharply with ϕ, and increases with ρ.

The marginal of BWE_P in ϕ deserves a word, since it may seem to contradict the damping story: producer bullwhip is *higher* at $\phi = 1.0$ than at $\phi = 0.6$. This is exactly the mechanism at work. A higher yield means the constraint binds less often, so less of the order tail is clipped, so the producer order variance is closer to the (higher) linear-chain value. Low yield binds often, clips heavily, and drives bullwhip down—at the cost of the high constraint frequency shown alongside. The two columns must be read together: the bullwhip reduction at low yield is bought with frequent binding, which is the operational penalty quantified in Section 5.6.

Table 11: Marginal means of producer bullwhip BWE_P and constraint frequency by factor level, averaged over all other factors (full 576×30 design per policy).

		OUT		POUT	
Factor	Level	BWE_P	Constr.	BWE_P	Constr.
α	0.2	60.4	0.201	27.5	0.377
	0.4	166.3	0.085	29.9	0.332
	0.6	262.9	0.055	33.9	0.286
	0.8	375.8	0.041	41.9	0.241
t	1	127.4	0.143	5.7	0.482
	2	214.9	0.084	33.4	0.263
	3	306.7	0.060	60.7	0.182
ρ	0.2	236.7	0.078	34.6	0.282
	0.3	224.0	0.088	33.7	0.299
	0.4	210.1	0.101	32.8	0.318
	0.5	194.7	0.116	32.0	0.338
τ	1	216.1	0.096	33.3	0.321
	3	216.3	0.095	33.3	0.312
	5	216.5	0.096	33.3	0.305
	10	216.5	0.096	33.2	0.300
ϕ	0.6	190.2	0.181	31.2	0.528
	0.8	211.2	0.099	32.4	0.393
	1.0	247.6	0.008	36.2	0.007