OMNITEK CONSULTING

Solving the IRT Waste Problem

WHITE PAPER

How Smarter Randomization and Trial Supply System Design Cut Drug Waste and Site Frustration

Abstract

Clinical trial investigational product waste sits at a median of roughly 50%, according to McKinsey research into clinical supply chains, meaning half the drug supply a sponsor manufactures, labels, and ships for a study is never actually dispensed to a patient. That waste is not primarily a manufacturing problem. It traces back to how randomization and trial supply management, the interactive response technology that decides which patient gets which kit and when a site gets resupplied, gets configured at study startup and then rarely revisited as enrollment reality diverges from the forecast built months earlier. Poorly tuned resupply thresholds create a second cost nobody puts on a budget line: site coordinators who spend real time chasing kits that should have already arrived, or who lose confidence in a system that keeps them guessing. IRT-enabled forecasting can cut supply chain waste by an estimated 15 to 20%, but only when the system is actively tuned against real enrollment and site-level consumption data rather than treated as a one-time build. This white paper applies established data-driven forecasting and predictive-signal best practices to IRT and RTSM optimization, giving clinical supply and operations leaders a practical approach to turning trial data into the resupply signal it should have been from the start.

The Waste Nobody Budgets For

Investigational product waste in clinical trials sits at a median of roughly 50%, according to McKinsey's research into clinical supply chain performance, a figure the firm attributes largely to forecasting and planning gaps rather than to manufacturing inefficiency itself. Half of what a sponsor spends producing, labeling, and shipping drug supply for a study never reaches a patient.

That waste hides in plain sight because it rarely shows up as a single line item. It accumulates across dozens of smaller decisions: a randomization scheme built around an enrollment curve that turned out to be optimistic, a resupply threshold set conservatively at study startup and never revisited, expiring kits at sites that enrolled slower than planned while other sites ran short. McKinsey's broader supply chain review also found that 19% of packaged and labeled kits show some form of deviation and 7% generate a formal complaint, numbers that point to a supply process under real operational strain well before a single patient is dosed.[i]

What IRT and RTSM Systems Are Actually Deciding

Interactive response technology, sometimes called randomization and trial supply management, is the system that executes two connected jobs during a study: assigning each enrolled patient to a treatment arm according to the protocol's randomization scheme, and triggering resupply shipments to sites based on current and projected inventory. Every kit a site holds, every threshold that triggers a new shipment, and every rule governing how substitutions or dose adjustments get handled all live inside this one system.

The system gets built once, during study startup, against a forecast of enrollment pace and site-level consumption that is, at that point, still a projection. Actual enrollment almost never tracks that projection precisely, some sites activate late, some enroll faster than expected, some see higher-than-planned dropout. The IRT configuration, if left as originally built, keeps operating against the original assumptions long after reality has diverged from them, which is precisely how conservative resupply thresholds turn into both surplus at slow sites and stockouts at fast ones inside the same trial.

Applying Best Practices: Turning Trial Data Into a Live Resupply Signal

A foundational best practice in clinical supply forecasting, treating operational data as a continuous signal to act on rather than a static input configured once and left alone, applies directly to IRT and RTSM optimization. Applied here, it means the resupply logic driving shipments to sites should update against real enrollment and consumption data throughout the trial, not just at the initial build.

1. **Scheduled Forecast Recalibration Against Actual Enrollment.** Set a defined cadence, not an ad hoc one, for comparing actual site-level enrollment and consumption against the original forecast, and adjusting resupply thresholds accordingly rather than waiting for a stockout or a visible surplus to force the conversation.
2. **Site-Level Consumption Data as the Real Trigger, Not a Fixed Buffer Alone.** Weight resupply timing toward each site's actual consumption pattern rather than relying solely on a uniform buffer built for the average site, since a single threshold applied across sites with very different enrollment paces is exactly what produces simultaneous surplus and shortage.

3. **IRT Configuration Change as a Managed, Not Avoided, Process.** Build a lightweight but real process for adjusting IRT configuration mid-trial, since the alternative, leaving an outdated configuration in place to avoid the change effort, is what generates most of the waste in the first place.
4. **Randomization Scheme Sensitivity Reviewed Alongside Supply Data.** Periodically confirm that the randomization scheme itself, block size, stratification factors, is still performing as intended against actual enrollment patterns, since a scheme drifting from its design assumptions can itself distort supply demand at individual sites.

None of this requires rebuilding the IRT system mid-study. It requires treating the data the system already collects as an ongoing signal worth acting on, rather than a record kept for the close-out report.

What a Live-Signal Approach Buys, and Where It Gets Hard

Advantages

- **Meaningful reduction in manufactured and shipped waste:** IRT-enabled forecasting tuned against real consumption data can cut supply chain waste by an estimated 15 to 20%, directly against a median baseline near 50%.
- **Fewer site-level stockouts and the frustration that follows:** resupply thresholds calibrated to actual site consumption reduce the specific failure mode that erodes site confidence in the sponsor's operational competence.
- **Better-informed manufacturing and comparator sourcing decisions:** a resupply signal grounded in real enrollment data feeds more accurate demand projections back into procurement rather than compounding the original forecast's errors.
- **Reduced expiry-driven waste at slower-enrolling sites:** recalibrated thresholds catch a site's actual pace early enough to adjust shipment volume before excess kits sit past their usable window.

Risks and Barriers

- **IRT changes require formal change control and can carry real cost:** mid-trial configuration adjustments touch a validated system and need to go through the same change management rigor as any other GxP system update.
- **Recalibration requires genuine cross-functional coordination:** connecting clinical operations enrollment data, site-level consumption data, and IRT configuration decisions is an organizational process as much as a technical one.
- **Comparator and cell and gene therapy supply chains add complexity:** sourcing constraints and short shelf life narrow the margin for error when tuning resupply thresholds for these categories specifically.
- **A scheduled recalibration cadence competes for attention with go-live pressure:** clinical operations teams focused on hitting enrollment milestones can deprioritize a periodic supply data review that has no single obvious owner.

Case in Point: The Trial That Ran Out at Its Busiest Sites

A mid-size sponsor running a multi-region Phase III trial built its IRT resupply logic at study startup around a projected enrollment curve that assumed roughly even activation across all 60 participating sites. Actual enrollment did not follow that pattern. A cluster of a dozen sites in one region activated early and enrolled well ahead of projection, while several sites in another region faced local IRB delays and enrolled several months behind schedule.

Because the resupply thresholds had been set once against the original average-pace forecast and never revisited, the fast-enrolling sites began running short of kits roughly four months into the trial, triggering emergency shipments and, at two sites, brief enrollment pauses while coordinators waited on replenishment. Meanwhile the slower-enrolling sites accumulated kits well beyond what their actual pace required, several of which expired before being dispensed. Site coordinators at the fast-enrolling sites reported the stockouts directly to the sponsor's clinical operations team as a recurring source of frustration that was affecting their willingness to prioritize the study among competing trials.

The sponsor's fix was establishing a monthly review comparing actual site-level consumption against the original forecast, adjusting resupply thresholds site by site rather than applying a single trial-wide buffer. Within two review cycles, the emergency shipment pattern at the fast-enrolling cluster stopped recurring, and the sponsor redirected a portion of the excess inventory accumulating at slower sites before it reached expiry.

Recommendations for Clinical Supply and Operations Leaders

1. **Establish a defined recalibration cadence for IRT resupply thresholds.** rather than treating the initial startup configuration as fixed for the study's duration.
2. **Give site-level consumption data a direct route into supply decisions.** so resupply logic reflects actual site pace rather than a single trial-wide average.
3. **Build a lightweight, pre-approved change control pathway for IRT configuration updates.** so recalibration does not require a lengthy ad hoc approval process each time it is needed.
4. **Assign clear ownership for the enrollment-to-supply feedback loop.** since a review that has no specific owner tends to lose out to more visible enrollment-milestone pressure.
5. **Apply extra recalibration attention to comparator and cell and gene therapy supply.** given the tighter sourcing constraints and shorter shelf life these categories carry.
6. **Track site-reported supply frustration directly, not just stockout counts.** since a site's confidence in the resupply process affects enrollment prioritization independent of whether a stockout technically occurred.

Measuring Success: KPIs for IRT and Trial Supply Optimization

- **Investigational product waste rate:** the percentage of manufactured and shipped kits never dispensed to a patient, tracked toward improvement against the roughly 50% industry median baseline.
- **Site-level stockout incidents:** the number of resupply shortfalls severe enough to affect dosing or enrollment continuity, tracked per site and per quarter.

- **Expiry-driven kit loss:** the volume of kits that expire unused at a site, an indicator of resupply thresholds set too high against actual site pace.
- **Recalibration cycle adherence:** the percentage of scheduled forecast-versus-actual reviews completed on cadence, tracked as a leading indicator for the other three metrics.

Conclusion

A median waste rate near 50% is not a manufacturing problem waiting for a better production process. It is a signal problem, an IRT configuration built once against a forecast and then left to operate on assumptions the trial has already outgrown. Sites feel the consequence directly, in stockouts and expired kits and the erosion of confidence that follows both.

Applying disciplined best practices to IRT and RTSM optimization, scheduled recalibration, site-level consumption as the real trigger, a managed pathway for configuration change, treats trial supply data as the live resupply signal it was always capable of being, closing a substantial share of the gap between what a sponsor manufactures and what a patient actually receives.

About OmniTek & Contact Information

OmniTek Consulting is a trusted partner for organizations seeking to enhance governance, operational agility, and decision-making. We specialize in modernizing business processes with a blend of human-centered design, technology expertise, and hands-on change management.

For clinical supply and operations leaders looking to reduce drug waste and site-level frustration, OmniTek applies proven best practices to build a practical IRT and RTSM optimization practice: scheduled recalibration against real enrollment data, site-level resupply logic, and a managed pathway for mid-trial configuration change. Whether you are scoping your next study's IRT build or diagnosing waste and stockout patterns in a trial already underway, OmniTek can help you turn your trial supply data into a signal that actually drives better decisions.

Contact us today to explore how OmniTek can support your transformation journey and drive lasting results.

McLean, VA | Serving clients nationwide
info@omnitekconsulting.com
omnitekconsulting.com

Click here to book a 15-minute call to see how OmniTek can help your team put these insights into practice.

Endnotes

[i] McKinsey & Company, "Clinical supply chains: How to boost excellence and innovation," on median investigational product waste levels and packaged-kit deviation and complaint rates, available at https://www.mckinsey.com/industries/life-sciences/our-insights/clinical-supply-chains-how-to-boost-excellence-and-innovation

[ii] Suvoda, "IRT: Driving supply and cost optimization," on estimated waste and cost reductions from IRT-enabled forecasting and industry survey data on supply cost drivers, available at https://www.suvoda.com/insights/blog/irt-driving-supply-cost-optimization