

WHITEPAPER

# AI on the Shop Floor

## Practical Use Cases for Midmarket Manufacturers

Beyond the hype: where AI is already delivering measurable results in quoting, scheduling, quality, and decision-making.

---

### ERP Outcomes Consulting

Outcomes-Driven ERP Selection for Manufacturers

[erpoutcomes.org](http://erpoutcomes.org)

## EXECUTIVE SUMMARY

## AI for Manufacturers: Separating Signal from Noise

---

Every ERP vendor, every consulting firm, and every conference keynote is talking about AI. Most of the conversation is abstract. Generative AI will transform manufacturing. Machine learning will revolutionize supply chains. Predictive analytics will change everything.

Manufacturers are left with a reasonable question: what does any of this actually mean for my shop floor, my quoting process, and my bottom line?

This paper cuts through the hype. It presents practical, specific AI use cases that are relevant to midmarket manufacturers today, not theoretical capabilities that require a data science team and a seven-figure investment, but concrete applications that can be implemented within or alongside modern ERP systems to produce measurable business outcomes.

Each use case is grounded in real manufacturing challenges: quoting bottlenecks, scheduling complexity, quality variability, margin erosion, and the institutional knowledge problem that every manufacturer with an aging workforce is quietly worried about.

The goal is not to convince you that AI is important. It is to show you where it is already working and how to evaluate whether it belongs in your ERP strategy.

## 01

## SECTION ONE

## Intelligent RFQ Triage and Quoting

---

For most midmarket manufacturers, quoting is the single highest-leverage process in the business. Every dollar of revenue starts as a quote, and yet, quoting is where some of the largest invisible losses hide.

### The Problem

A precision machining company receives 25 or more RFQs per week. The quoting process requires coordination across engineering, estimating, and sales through a series of meetings and handoffs that take roughly 1.5 weeks to produce a confident quote. The result: the team can realistically focus on one quote per week. The other 20 or more RFQs quietly expire without a response.

Nobody tracks those expired RFQs as lost revenue. There is no report, no alarm, no chargeback. The loss is completely invisible.

### Where AI Changes the Equation

**Automated RFQ triage.** An AI model trained on your historical quoting data, such as win rates, margins, part families, customer profiles, and material costs, can score incoming RFQs within minutes of receipt. Instead of the sales team triaging by gut feel, the system flags the highest-probability, highest-margin opportunities and deprioritizes the ones that are unlikely to convert or would produce unacceptable margins.

**Template-based estimating.** AI can match an incoming RFQ to the most similar historical quotes in your system, based on part geometry, material, tolerances, operations, and customer requirements—and generate a starting estimate automatically. The estimator refines rather than builds from scratch, compressing hours of work into minutes.

**Real-time cost accuracy.** Machine metrics fed into the quoting process replace assumptions with actual data. Instead of estimating cycle times from a reference book, the system pulls actual cycle times from your CNC machines for similar parts. Instead of using standard material rates, it uses current purchase prices. The quote reflects reality, not a best guess from last year.

### The Outcome

A manufacturer that implements this approach can realistically move from responding to 1 RFQ per week to 5 or more per week, targeting only the highest probability opportunities. The quoting cycle compresses from 1.5 weeks to hours. The win rate improves because the team is spending their time on the right opportunities with accurate pricing.

# 02

## SECTION TWO

### Predictive Production Scheduling

---

Production scheduling in most midmarket manufacturers is a combination of ERP-generated suggestions, tribal knowledge, and daily firefighting. The scheduler knows that machine 7 runs slow on Tuesdays because of the operator, that certain tooling combinations cause problems, and that Customer A's orders always need to ship a day early because their receiving dock closes at noon.

None of that knowledge lives in the system. It lives in one person's head.

#### Where AI Changes the Equation

**Pattern recognition from historical data.** AI can analyze years of production history to identify patterns that humans cannot see at scale: which job combinations cause bottlenecks, which sequences minimize changeover time, which machine-operator pairings produce the best throughput, and which orders are most likely to be expedited based on customer behavior.

**Predictive disruption modeling.** Instead of reacting to machine downtime, material shortages, and quality holds after they happen, AI models can predict them. A machine that has been gradually trending toward a maintenance threshold triggers a scheduling adjustment before the breakdown occurs. A supplier whose lead times have been creeping up is flagged before a stockout happens.

**Dynamic rescheduling.** When disruptions do occur—and they always do—AI can regenerate the schedule in seconds, optimizing across the full constraint set: machine availability, tooling, labor, material, and customer priority. A process that takes a human scheduler hours of phone calls and whiteboard revisions happens automatically.

#### The Outcome

Manufacturers using AI-assisted scheduling typically see on-time delivery improvements of 10 to 20 percentage points, a reduction in expediting costs, and a significant decrease in the organizational chaos that accompanies daily schedule changes. More importantly, the tribal knowledge that currently lives in one scheduler's head becomes embedded in the system—reducing a business risk that every manufacturer quietly worries about.

## 03

## SECTION THREE

## Quality Prediction and Connected Process Control

---

Traditional quality management is reactive. You produce the part, inspect it, find the defect, and either scrap it or rework it. The cost of quality—scrap, rework, customer chargebacks, warranty claims—is one of the most significant margin drains in manufacturing. For IATF-certified or aerospace manufacturers, the stakes are even higher.

### Where AI Changes the Equation

**In-process quality prediction.** Sensors on CNC machines, die cast presses, and welding stations generate continuous data streams: temperatures, pressures, vibration, cycle times, spindle loads. AI models trained on this data can predict when a process is drifting toward an out-of-spec condition before it produces a defective part. The operator receives an alert. The machine adjusts. The defect never happens.

**Root cause pattern analysis.** When defects do occur, AI can correlate them across dozens of variables simultaneously: which machine, which operator, which shift, which material lot, which tooling, which ambient temperature. Patterns that would take a quality engineer weeks of investigation to find surface in minutes.

**Supplier quality scoring.** AI can continuously evaluate incoming material quality against historical performance, flagging lots that are statistically likely to cause downstream problems before they enter production. This is particularly valuable in die casting and machining where material variability directly impacts yield.

### The Outcome

Connected process control powered by AI can reduce scrap rates by 20 to 40 percent, virtually eliminate customer-facing quality escapes, and transform the quality function from a cost center focused on inspection to a strategic capability focused on prevention. For IATF and aerospace manufacturers, this is the difference between audit findings and audit excellence.

## 04

## SECTION FOUR

## Real-Time Margin Intelligence

---

Most midmarket manufacturers cannot tell you what their actual margin is on a job until weeks after it ships. Estimated costs go into the quote. Actual costs trickle in over days or weeks as labor is posted, material is consumed, and outside processing invoices arrive. By the time the picture is complete, the opportunity to act on it has passed.

### Where AI Changes the Equation

**Live cost-to-date tracking.** AI-powered dashboards can aggregate real-time data from shop floor transactions, material consumption, labor reporting, and machine metrics to show actual cost versus estimated cost at any point during production—not after the fact.

**Margin erosion alerting.** When a job's actual costs begin trending above the estimate, the system alerts the production manager and the estimator in real time. The production manager can investigate and correct the root cause. The estimator can adjust future quotes for similar jobs. The feedback loop closes in hours, not months.

**Pricing optimization.** AI can analyze historical job profitability across part families, customers, volumes, and material types to identify systematic pricing inaccuracies. If your standard-rate-per-hour assumption underestimates the true cost of complex setups by 15 percent, the system will surface that pattern—along with the cumulative margin impact.

### The Outcome

Manufacturers who implement real-time margin visibility typically recover 2 to 5 margin points within the first year—not through cost cutting, but through pricing accuracy, early intervention on troubled jobs, and the elimination of systematic estimating errors that had been silently eroding profitability for years.

*The margin you think you're earning and the margin you're actually earning are rarely the same number. AI makes the gap visible in time to do something about it.*

# 05

## SECTION FIVE

# Institutional Knowledge Capture and the Workforce Transition

Every midmarket manufacturer has the same quiet fear: what happens when the people who know how everything works retire?

The scheduler who knows which jobs to prioritize. The estimator who knows how to price complex machined assemblies. The quality manager who knows which process parameters actually matter. The buyer who knows which suppliers will deliver and which ones won't. This knowledge is the operating system of the business, and it lives almost entirely in people's heads.

## Where AI Changes the Equation

**Decision pattern extraction.** AI can observe and codify the decision patterns of experienced operators, schedulers, and estimators. When your best scheduler consistently overrides the system's suggestion for certain job types, AI can learn why and begin incorporating that logic. The tribal knowledge doesn't disappear when the person retires, it becomes embedded in the system.

**Intelligent work instructions.** AI-powered work instruction systems can adapt in real time based on the operator's experience level, the specific part being produced, and historical quality data. A veteran machinist sees a streamlined instruction set. A new hire sees step-by-step guidance with visual aids and the specific parameters that matter most for quality on this particular operation.

**Conversational system access.** Natural language interfaces allow shop floor workers to interact with the ERP by asking questions in plain language: "What's the status of order 4523?" "When is the next delivery from supplier X?" "Show me the routing for part number 1847." This lowers the barrier to entry for new hires and reduces training time.

## The Outcome

AI does not replace experienced manufacturing professionals. It captures and scales their expertise so that the next generation can be productive faster, the business is less vulnerable to turnover, and institutional knowledge becomes a permanent organizational asset rather than a liability that walks out the door every time someone retires.

# 06

## SECTION SIX

# Demand Sensing and Intelligent Procurement

---

Midmarket manufacturers typically manage inventory through a combination of MRP calculations, safety stock buffers, and buyer intuition. When supply chains are stable, this works well enough. When they are volatile, as they have been for the past several years, it produces either stockouts that halt production or excess inventory that ties up cash.

## Where AI Changes the Equation

**Demand pattern recognition.** AI can analyze order history, customer behavior, seasonality, and external signals (commodity prices, economic indicators, industry data) to generate demand forecasts that are significantly more accurate than traditional MRP calculations based on historical averages.

**Dynamic safety stock optimization.** Instead of static safety stock levels set once a year, AI continuously adjusts buffer quantities based on current demand volatility, supplier reliability, and lead time variability. Items with reliable supply get leaner buffers. Items with volatile supply get more protection. The system adapts in real time.

**Supplier risk monitoring.** AI can track supplier performance trends, such as delivery accuracy, quality trends, and lead time changes, flagging deteriorating performance before it causes a production impact. This is especially valuable for single-source components where a disruption has outsized consequences.

## The Outcome

Manufacturers using AI-driven procurement typically see a 15 to 25 percent reduction in inventory carrying costs alongside a reduction in stockouts. The buyer's role shifts from transactional purchasing to strategic supplier management, which is the kind of work that actually moves the needle on cost and reliability.

## 07

## SECTION SEVEN

## Where to Start: An Outcomes-Driven Approach to AI

---

The biggest risk with AI in manufacturing is not that it won't work. It's that companies invest in the wrong use case first, don't see results, and conclude that AI isn't relevant to their business.

The right starting point is not the most technically impressive use case. It is the one that connects directly to a measurable business outcome your organization is already trying to achieve.

### Start with the Outcome, Not the Technology

If your biggest challenge is quoting speed, start with intelligent RFQ triage. If your biggest challenge is on-time delivery, start with predictive scheduling. If your biggest challenge is margin erosion, start with real-time job costing. Let the business outcome select the AI use case, not the other way around.

### Evaluate AI Capabilities During ERP Selection

If your organization is evaluating ERP systems, AI capabilities should be part of the evaluation, but evaluated through the outcomes lens, not as a feature checkbox. The right question is not "Does your platform have AI?" It is "Show me how your AI capabilities would help us achieve 24-hour quote turnaround for our top part families."

### Don't Boil the Ocean

The manufacturers who succeed with AI start with one use case, prove the value, and expand. They don't try to implement predictive scheduling, quality prediction, demand sensing, and margin intelligence all at once. They pick the use case with the clearest business case, implement it well, and let the results build organizational confidence for the next one.

### Ensure Your Data Foundation Is Solid

AI is only as good as the data it learns from. If your ERP has years of inaccurate inventory records, incomplete job costing, or inconsistent labor reporting, the AI will learn from bad data and produce bad recommendations. A clean, well-maintained ERP with disciplined data entry is the prerequisite for every AI use case in this paper.

*AI doesn't require perfection. It requires a foundation. If your ERP data is reasonably clean and your processes are reasonably consistent, you have enough to start.*

## 08

## SECTION EIGHT

## AI Use Cases at a Glance

Use Case	Business Outcome	Typical Impact
<b>RFQ Triage &amp; Quoting</b>	Faster quote response, higher win rate	Quote cycle from weeks to hours; 3–5x more responses
<b>Predictive Scheduling</b>	Improved on-time delivery	10–20 point OTD improvement; reduced expediting
<b>Quality Prediction</b>	Lower cost of quality, fewer escapes	20–40% scrap reduction; near-zero customer defects
<b>Margin Intelligence</b>	Pricing accuracy, margin recovery	2–5 margin points recovered in year one
<b>Knowledge Capture</b>	Workforce resilience, faster onboarding	30–50% reduction in new hire ramp time
<b>Demand &amp; Procurement</b>	Inventory optimization, fewer stockouts	15–25% reduction in carrying costs

### Start the Conversation

ERP Outcomes Consulting helps midmarket manufacturers evaluate ERP systems through an outcomes-driven methodology, including the AI capabilities that matter for your specific business challenges.

If you're evaluating ERP and want to understand how AI fits into your strategy, let's talk.

### ERP Outcomes Consulting

info@erpoutcomes.org | 763-900-9827 | erpoutcomes.org

*ERP Outcomes Consulting is vendor-neutral. We do not resell software or receive referral fees from ERP vendors. Our only commitment is to helping manufacturers make the best possible technology decision for their business.*

© 2026 ERP Outcomes Consulting, a Roberts Consulting Group LLC company. All rights reserved.