Cost–Quality–Energy Optimization for a Smart Manufacturing Line Using a Public Industrial Metrics Dataset

Tasbirul Alam Zihan [a]
[a] School of Industrial and Systems Engineering, University of Oklahoma

Abstract

Modern smart-manufacturing lines face tight energy budgets and stringent quality targets while delivering stable, low-cost output. This study presents a transparent, operations-ready optimization framework: public shop-floor data (production, maintenance, energy, downtime, defects) are distilled into conservative, auditable guardrails via medians and quantiles, and an interpretable linear program selects daily production, maintenance hours, and downtime under time, energy, and quality limits. The model is solved with Gurobi and interpreted through shadow prices, one-at-a-time sensitivity, and ε-constraint Pareto frontiers for cost–energy and cost–quality trade-offs. The baseline indicates demand, rather than the capacity, governs the operating point (binding demand, marginal maintenance minimum, slack time/energy/downtime). Sensitivity and frontier analyses show costs remain flat under moderate tightening of energy or quality guardrails and rise sharply only near the energy feasibility wall or with substantially stricter quality. The outcome is a simple, reproducible planning tool that quantifies the value of extra demand, time, energy, or quality and supports day-to-day scheduling.

Keywords: Production planning; Linear programming; Guardrails; Energy–quality trade-offs.

1 Introduction

Manufacturing planners must meet daily output while staying within energy budgets and maintaining quality. This paper presents a transparent optimization framework that turns routine shop-floor records into conservative, auditable guardrails on time, energy, downtime, and minimum maintenance, and then solves an interpretable linear program to recommend a single-day plan at minimum cost. A public Kaggle dataset serves as the case study. After light cleaning, seven operational fields are retained because they align with controllable levers and common plant limits: production volume, production cost, maintenance hours, energy consumption, additive process time, downtime percentage, and defect rate. Robust summaries (medians and tail percentiles) are used to set planning caps and coefficients so every number can be traced directly to the data. The model is solved with Gurobi and evaluated with three standard lenses: baseline feasibility and sanity checks, one-at-a-time sensitivity with shadow-price interpretation, and Pareto frontiers for cost versus energy and cost versus maintenance-driven quality. The contribution is a simple, reproducible data-to-decision pipeline that yields an explainable plan, clear trade-off suggestions, and manager-friendly discussions.

2 Literature Review

Production planning for modern manufacturing balances multiple, often conflicting objectives: low operating cost, reliable throughput, stringent quality, and tight energy budgets. Energy-aware production and scheduling has been studied from operational strategies to full optimization models that embed energy costs and caps alongside classical metrics.

In this line, Mouzon et al. (2007) showed that switching underutilized machines OFF during sufficiently long idle windows yields measurable energy savings with only small increases in completion time, motivating energy-aware dispatching and mixed-integer formulations. Complementarily, Liu et al. (2014) proposed a bi-objective job-shop model that minimizes total electricity consumption and weighted tardiness, using NSGA-II to reveal the energy–tardiness trade-off. In some studies, quality, maintenance, and planning are tightly coupled. Rivera-Gómez et al. (2018) developed an integrated production–maintenance model for a deteriorating machine where failure behavior and product quality depend on the production rate; the optimal policy alternates between minimal repair and major maintenance to minimize long-run cost.

When objectives conflict, the ε-constraint method remains a transparent way to map trade-offs by treating all but one objective as hard caps and solving a sequence of single-objective programs Haimes (1971). On the data side, regression quantiles Koenker and Bassett Jr (1978) motivate median/quantile summaries as outlier-resistant, operationally defensible statistics for parameterizing guardbands. These strands motivate our use of quantile-based summaries to set conservative guardrails linking production, time, energy, maintenance, and downtime, followed by a compact optimization model to reveal cost–energy–quality trade-offs and interpret them via dual analysis and ε-constraint frontiers.

3 Problem Description and Methodology

A single production day for one manufacturing line is considered. The line must meet a demand target while respecting limits on available hours (time capacity), daily energy, and quality. Quality is enforced indirectly through a minimum level of planned preventive maintenance and a cap on allowable downtime. The problem is formulated as a linear program (LP) and solved to optimality. No index sets are required because a single line and a single day are modeled. Parameters are estimated from the dataset described in Section 4.

The model selects three nonnegative quantities as decision variables for each day, which are,

The production volume x (units/day),

The planned preventive maintenance time m (hours/day), and

The downtime fraction $u \in [0, 1]$ representing the share of the day the line is unavailable.

Definitions and estimation details of parameters are listed in Table 1

Table 1: Parameters (estimated from data via robust quantiles)

Parameter	Meaning (units)
α	Hours per unit (h/unit), median of AdditiveProcessTime/ProductionVolume
g_1	Energy per unit (kWh/unit), 90th percentile of EnergyConsumption/ProductionVolume
g_2	Energy per maintenance hour (kWh/h), 90th percentile of positive residual after $g_1 x$
$H_{\max}$	Time capacity (h/day), 95th percentile of AdditiveProcessTime + small buffer, ≤ 24
$E_{\max}$	Energy cap (kWh/day), 95th percentile of EnergyConsumption
$m_{\min}$	Minimum maintenance (h/day), 25th percentile on low-defect days
$u_{\cap}$	Downtime ceiling (fraction), 90th percentile on low-defect days (truncated at 0.05)
$D_{\min}$	Demand floor (units/day), median of ProductionVolume
c_{unit}	Unit production cost (\$/unit), median of ProductionCost/ProductionVolume
c_{maint}	Maintenance cost (\$/h), median residual cost per MaintenanceHours
p_e	Energy price (\$/kWh)

The auxiliary energy variable can be eliminated by substitution. Because quantile-based caps can be mutually tight, a one-time, minimal enlargement of $H_{\max}$ and/or $E_{\max}$ is computed so that producing $D_{\min}$ at $m_{\min}$ is feasible without artificial slacks. With a small buffer $\beta \in [0.01, 0.02]$:

$$H_{\text{need}} = \alpha D_{\min} + m_{\min} \qquad H'_{\max} = \max\{H_{\max}, (1+\beta)\, H_{\text{need}}\} \tag{1}$$

$$E_{\text{need}} = g_1 D_{\min} + g_2 m_{\min} \qquad E'_{\max} = \max\{E_{\max}, (1+\beta)\, E_{\text{need}}\} \tag{2}$$

The baseline LP is then solved with $H'_{\max}, E'_{\max}$. This preserves the demand target and yields clean dual (shadow) prices. The baseline model is written as,

$$\min_{x,m,u} \quad c_{\text{unit}}\, x \;+\; c_{\text{maint}}\, m \;+\; p_e\,(g_1 x + g_2 m) \tag{3}$$

$$\text{s.t.} \quad \underbrace{\alpha\, x + m}_{\text{time use}} \;\le\; H'_{\max} \qquad \text{(time capacity)} \tag{4}$$

$$\underbrace{g_1 x + g_2 m}_{\text{energy use}} \;\le\; E'_{\max} \qquad \text{(energy cap)} \tag{5}$$

$$u \;\le\; u_{\cap} \qquad \text{(downtime guardrail)} \tag{6}$$

$$m \;\ge\; m_{\min} \qquad \text{(maintenance minimum)} \tag{7}$$

$$x \;\ge\; D_{\min} \qquad \text{(demand floor)} \tag{8}$$

$$x,\, m,\, u \;\ge\; 0 \tag{9}$$

Constraint (4) aggregates productive time and planned preventive maintenance. Constraint (5) links energy to production and maintenance activity using conservative, data-driven coefficients. Constraints (6)–(8) implement historical quality guardrails and the daily demand requirement. The objective (3) is fully interpretable as the sum of unit, maintenance, and purchased-energy costs.

The LP is solved to optimality in Gurobi, and we report the baseline plan $(x^\star, m^\star, u^\star)$ with implied energy, dual analysis with shadow prices for the main constraints (and reduced costs), a $\pm 10\%$ one-at-a-time sensitivity on $D_{\min}$, $H_{\max}$, and $E_{\max}$, and ε-constraint Pareto frontiers for cost–energy and cost–quality. Together, these outputs quantify trade-offs among cost, time, energy, and quality and provide an auditable basis for day-to-day planning.

4 Data Description

A publicly available manufacturing dataset from Kaggle (https://www.kaggle.com/datasets/rabieelkharoua/predicting-manufacturing-defects-dataset) is used as the empirical basis. Each record corresponds to one operating day of a single production line, and the cleaned working file is built. The analysis uses the following columns: ProductionVolume (units/day), ProductionCost (\$/day), MaintenanceHours (hours/day), EnergyConsumption (kWh/day), AdditiveProcessTime (hours/day), DowntimePercentage (fraction of day), and DefectRate (defects per 1k units). All fields were coerced to numeric types and rows missing required values were dropped. DowntimePercentage was expressed as a $[0, 1]$ fraction when provided in percent and simple sanity checks for non-negative hours and energy were applied. No imputation or model-based filtering was performed.

The historical panel is used only to compute robust, auditable guardrails and cost coefficients summarized in Tables 1. The linear program in Section 3 is then solved for a single representative day using these data-derived parameters.

5 Results and Discussion

Complete tables and figures that represent all the results (baseline solution, duals, sensitivity analyses, and ε-constraint frontiers) are deferred to the Appendix, and only the interpretation of results and insights is discussed here. The baseline LP solves to optimality and recommends producing about 549 units/day with roughly 6 hours of preventive maintenance and zero planned downtime; implied energy use is about 9×10^3 kWh/day and total cost is \$$1.35 \times 10^4$/day. Duals indicate that *demand* is the active driver of cost (positive shadow price), whereas time and energy caps are locally slack (zero shadow prices). The maintenance minimum is weakly binding with a very small positive dual, implying that modest increases in preventive maintenance raise cost only marginally.

Single-factor perturbations corroborate what the duals indicate. Tightening the time cap quickly makes time the bottleneck, while tightening the energy cap has no effect until it drops below the baseline energy requirement, at which point the problem becomes infeasible under current guardrails. By contrast, raising the demand target increases cost roughly at the demand shadow price, and lowering it reduces cost proportionally. Pareto frontiers confirm these patterns: the cost–energy curve is flat above baseline usage and breaks at the feasibility wall, whereas the cost–quality (via maintenance) curve is nearly linear with a shallow slope. Managerially, this suggests three actions: (i) cost is demand-limited at baseline, so gains come from lowering per-unit production/energy terms rather than adding capacity; (ii) energy policy should target process improvements that reduce per-unit kWh, not just tighter daily caps; and (iii) increasing preventive maintenance is a low-cost lever to support quality or uptime objectives.

6 Conclusion

This study presents a transparent, data-driven linear program for daily production planning under time, energy, and quality guardrails derived from robust statistics. By using medians and quantiles, every coefficient is auditable, and the solution is interpretable via shadow prices and feasibility checks. The baseline indicates that demand rather than the capacity governs the operating point: the demand constraint binds, the maintenance minimum is marginal, and time, energy, and downtime limits are slack. Sensitivity and Pareto analyses reinforce insights, including that the cost–energy curve is flat until a feasibility wall at baseline usage (implying true savings require lower per-unit kWh, not tighter caps), while the cost–quality curve is nearly linear with a shallow slope, consistent with a low marginal cost of extra maintenance.

The framework offers a practical decision aid that helps managers to quantify the value of an extra unit of demand, an hour of capacity, or a kilowatt-hour, and defend plans with simple tables and duals. Limitations include a single line, single-day horizon, linear/stationary coefficients, exogenous demand, and quality enforced via guardrails. Future work will extend to multi-period and multi-line settings, time of use energy pricing, robust variants, and calibrated piecewise linear links that maintain interpretability while expanding feasibility.

References

Haimes, Y. (1971). On a bicriterion formulation of the problems of integrated system identification and system optimization. *IEEE transactions on systems, man, and cybernetics* (3), 296–297.

Koenker, R. and G. Bassett Jr (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.

Liu, Y., H. Dong, N. Lohse, S. Petrovic, and N. Gindy (2014). An investigation into minimising total energy consumption and total weighted tardiness in job shops. *Journal of Cleaner Production 65*, 87–96.

Mouzon, G., M. B. Yildirim, and J. Twomey (2007). Operational methods for minimization of energy consumption of manufacturing equipment. *International Journal of production research 45*(18-19), 4247–4271.

Rivera-Gómez, H., O. Montaño-Arango, J. R. Corona-Armenta, J. Garnica-González, E. S. Hernández-Gress, and I. Barragán-Vite (2018). Production and maintenance planning for a deteriorating system with operation-dependent defectives. *Applied Sciences 8*(2), 165.