Towards Smarter Haulage Operations in an Open-Pit Limestone Mine

Quantifying the Impact of Dynamic Dispatching and a Single Autonomous Unit Using a Data-Driven DES

AESM7000: Master Thesis Applied Earth Sciences

Maurits Vermeer

Delft University of Technology

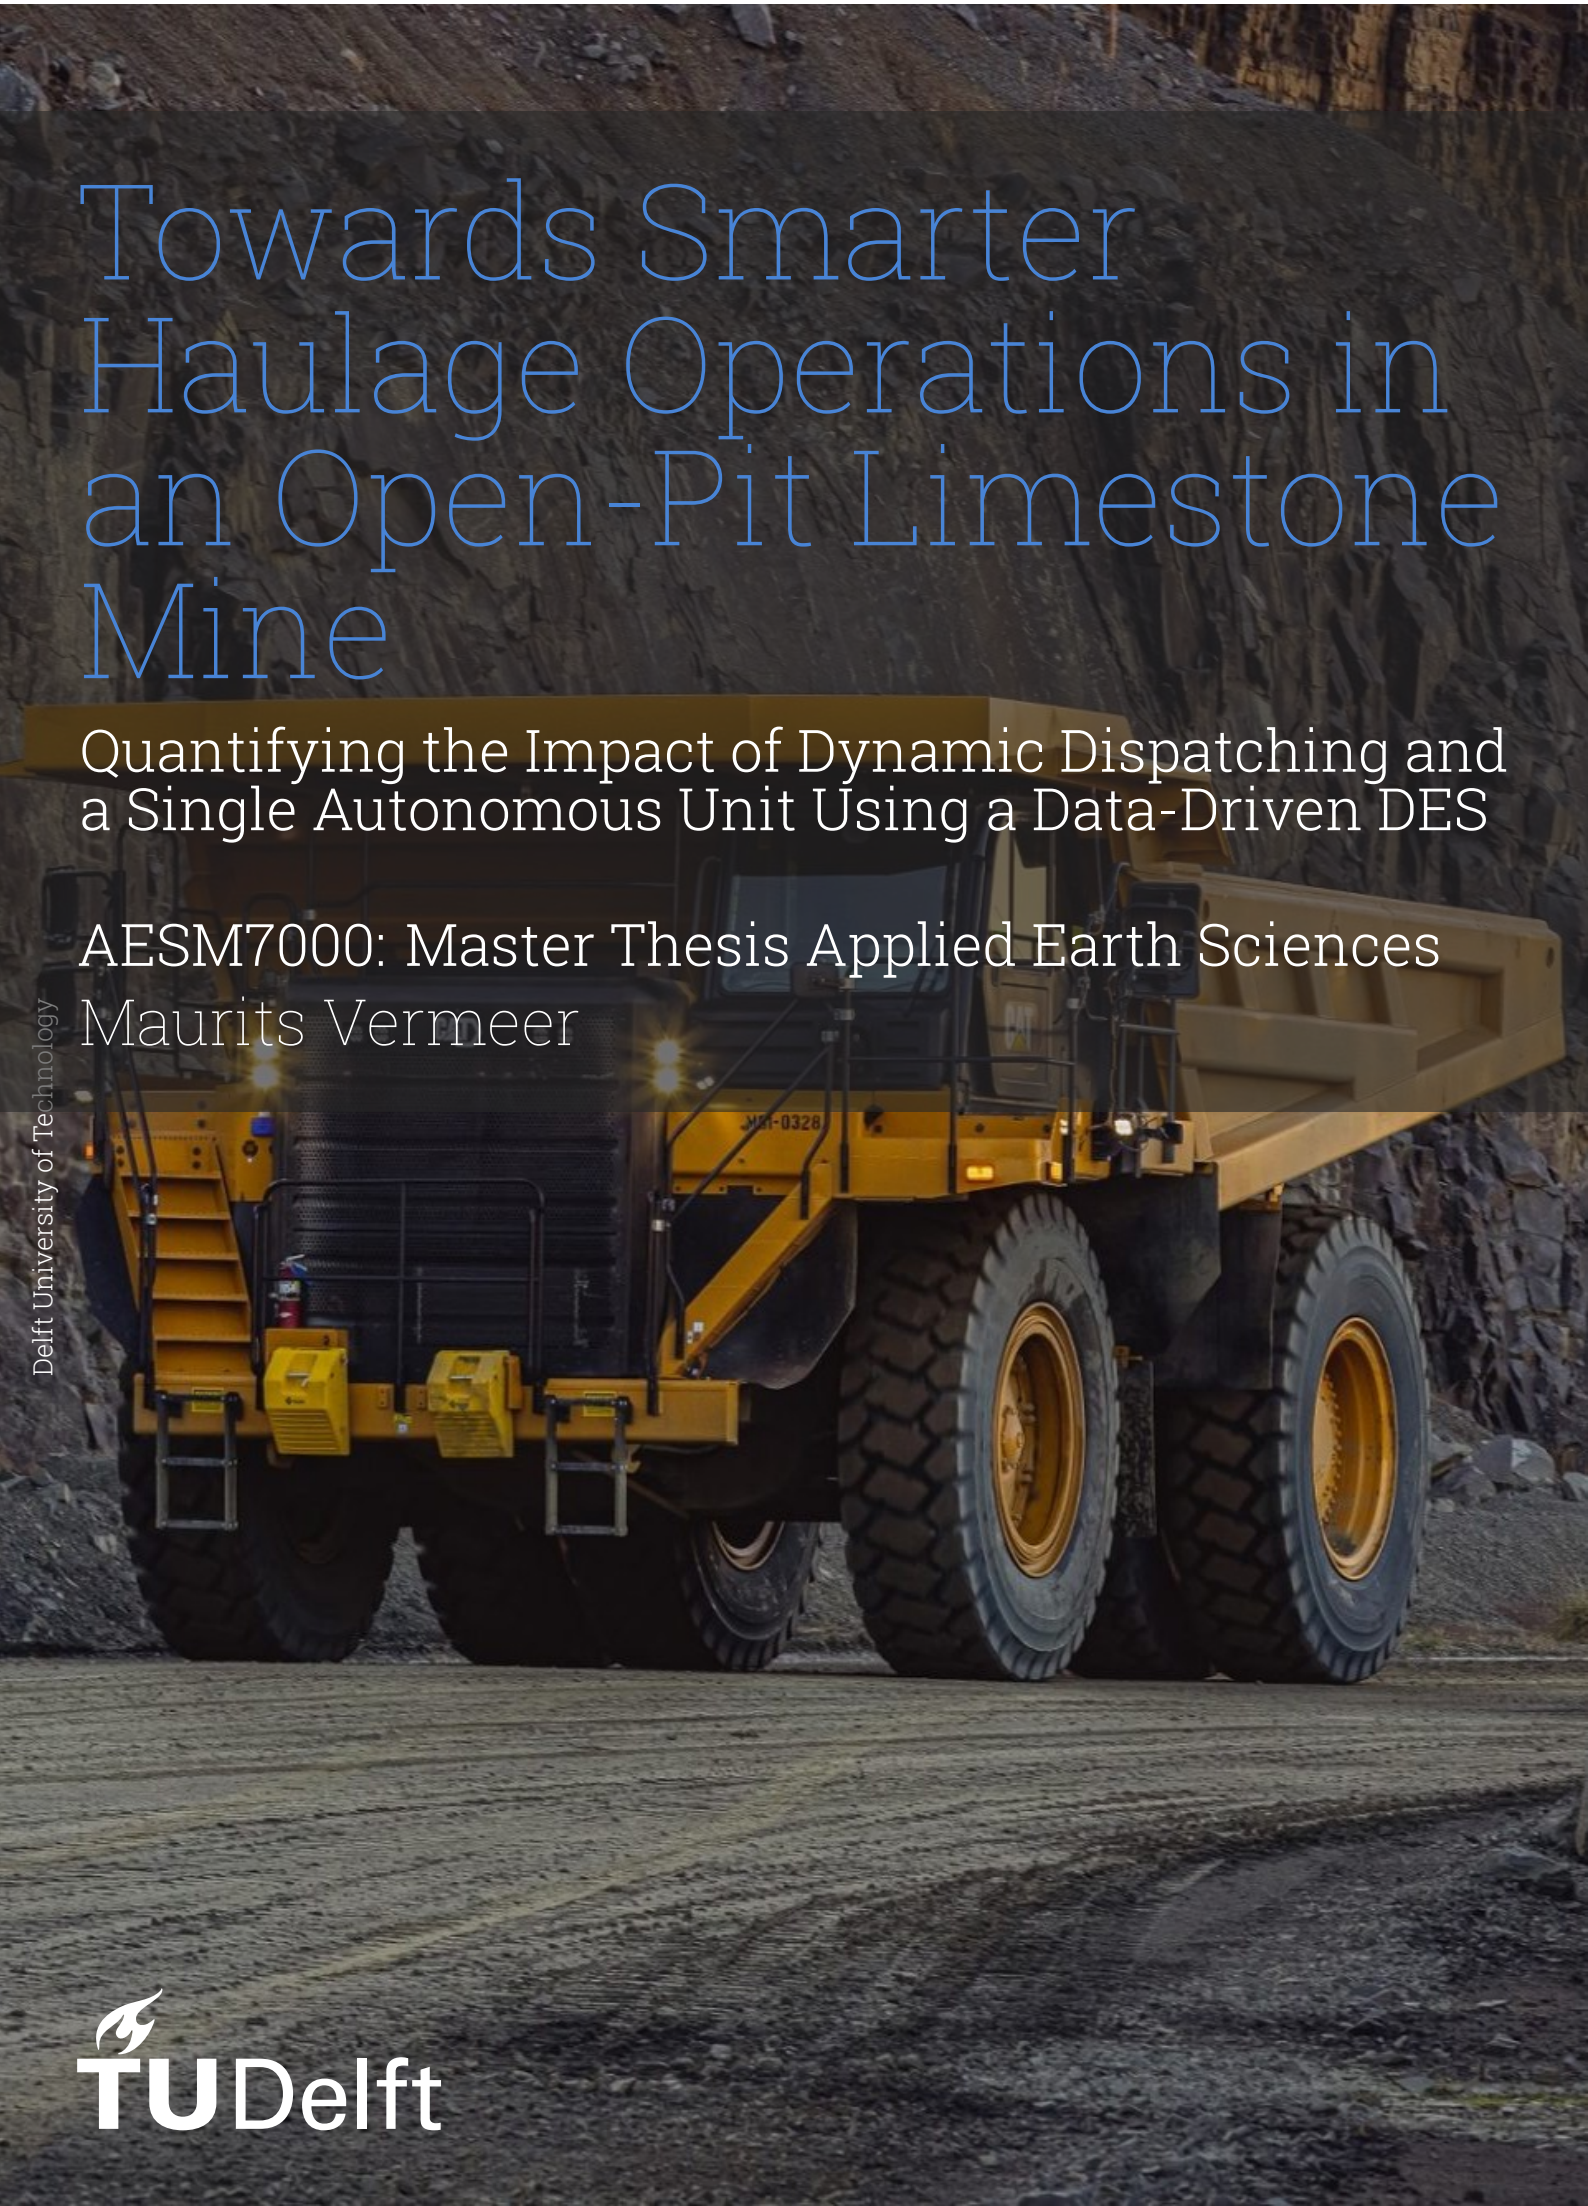

TU Delft

Towards Smarter Haulage Operations in an Open-Pit Limestone Mine

Quantifying the Impact of Dynamic Dispatching and a Single Autonomous Unit Using a Data-Driven DES

by

Maurits Vermeer

4964624

Instructors:	Dr. M. Buxton, Dr. F. Desta, Ir. M. Keersemaker, Dr. M. Ramgraber
Teaching Assistant:	K. Pashna
Project Duration:	March, 2025 - March, 2026
Faculty:	Faculty of Civil Engineering and Geosciences, Technical University Delft

Cover:	Autonomous Caterpillar 777 at Luck Stone limestone mine (Stewart, 2022)
Style:	TU Delft Report Style, with modifications by Daan Zwaneveld

TU Delft

Preface

This thesis marks the end of my Master's in Applied Earth Sciences at TU Delft. Spending a year working between the university and the Flandersbach mine in Germany was a great experience, and being on-site made the research far more practical than I initially expected.

I want to thank the team at the mine first. Dorian, the mine manager, gave me the chance to do my research there and I really appreciate that opportunity. A special thanks goes to Nico, who drove me to my office in the mine every single day. Both he and Dorian taught me how the industry actually works and were always there to help when I had questions.

At the university, Kiarash was extremely helpful. He attended every meeting and gave me the feedback I needed to keep the project on track. I also want to thank Mike for his guidance in our meetings. He helped me understand the industry perspective and taught me how to structure my research by looking for the operational bottlenecks first.

I am also grateful to Feven and Marco for reviewing my work. Marco's visit to the mine was a great addition to the process. Finally, I want to thank Max for taking the time in his busy schedule to be my last reviewer. I really appreciate him stepping in, especially after it proved difficult to find a final committee member.

On a personal note, I have to thank my parents. They helped me with the move to Germany and back, and it was great to have them visit the site so we could see the mine together. Lastly, I want to thank my grandfather. He is the one who originally got me interested in engineering. He was always available for a call and stayed genuinely interested in everything I was doing in Germany.

Maurits Vermeer
Delft, March 2026

Summary

This thesis quantifies the operational impact of two interventions at the Flandersbach limestone mine: a tuned heuristic dynamic dispatch strategy and autonomy-related fleet changes. A discrete event simulation digital twin was developed and calibrated using telemetry-derived empirical inputs and site-confirmed operating rules, and scenarios were evaluated over two representative operating weeks. The model is used to compare throughput, waiting and congestion behavior, and target adherence under consistent assumptions.

The results show that the implemented dynamic dispatch policy is not an effective mechanism to increase production volume in its current form. While the compliance metrics indicate that dynamic dispatch can enforce relatively tight adherence to shift-level bench targets under the thesis definition, it does so at a throughput cost. Compared to the fixed-assignment baseline, dynamic dispatch reduces total moved tonnage by about one percent and increases congestion, most notably through substantially higher loader queue time and headway-related waiting.

Single-agent autonomy produces asymmetric benefits across subsystems. Introducing a single autonomous hauler or loader yields limited gains in a mixed fleet because productive time remains bounded by shared schedules and downstream constraints. In contrast, the autonomous load-and-carry scenario produces the largest throughput increase in the core scenario set, nearly eight percent, consistent with the fact that this intervention expands the effective operating window by operating through periods that are otherwise constrained by breaks and shift transitions. This result should be interpreted in light of scope, since the load-and-carry unit is modeled as a dedicated crusher-feeding unit, whereas in practice it also performs auxiliary duties.

Overall, the findings indicate that improving a single component often yields diminishing system-level returns when the haulage cycle remains constrained by shared resources, schedules, and interaction effects. The discrete event simulation digital twin provides a validated sandbox for relative comparison of dispatch and autonomy concepts, and the results motivate future concepts that change system-level constraints more directly.

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
AHS	Autonomous Haulage System
CNP	Contract Net Protocol
CRN	Common Random Numbers
DD	Dynamic dispatch
DDQN	Double Deep Q-Network
DES	Discrete event simulation
KPI	Key Performance Indicator
LC	Load and carry
LHD	Load-Haul-Dump
MAS	Multi-Agent System
MP	Mathematical Programming
RL	Reinforcement Learning
SBO	Simulation-Based Optimization
TPE	Tree-structured Parzen Estimator

Symbols

Symbol	Definition	Unit
A	Intercept parameter in the grade–speed relationship	[-]
$\mathbf{AF}$	Aggressiveness factor	[-]
$\mathbf{AF}(t)$	Aggressiveness factor at simulation time step t	[-]
$\mathbf{AF}_{\min}$	Lower bound for the aggressiveness factor	[-]
$\mathbf{AF}_{\max}$	Upper bound for the aggressiveness factor	[-]
B	Grade sensitivity coefficient in the grade–speed relationship	[$\%^{-1}$]
$\text{Capacity}_{\text{Edge}}$	Maximum number of haul trucks allowed concurrently on an edge	[-]
$F_{\text{Objective}}$	Scalar objective value minimized during DD tuning	[t]
k	Variance factor used in the variance transformation	[-]
κ	Scaling parameter controlling the magnitude of behavioural variability	[-]
Length_{m}	Edge length in meters used in the safety-distance capacity proxy	[m]
P_{downtime}	Downtime penalty term reflecting crusher availability status	[-]
$P_{\text{empirical}}$	Value drawn from the empirical (manual) distribution	[varies]
P_{sample}	Transformed value used in the autonomous scenario	[varies]
P_{target}	Target-compliance signal used in DD scoring	[%]
T_{starve}	Loader starvation time used in DD scoring	[min]
T_{travel}	Travel-time estimate to a candidate loader used in DD scoring	[min]
T_{wait}	Queueing-delay proxy used in DD scoring	[s]
$Tonnage_{\text{Hauler}}$	Total hauled tonnage over the DD evaluation horizon	[t]

Symbol	Definition	Unit
v_{drive}	Realized driving speed after behavioural variability	[km/h]
$v_{\text{flat,pit,state}}$	Baseline speed on nominally flat grade for a pit and load state	[km/h]
$v_{\text{pit,state}}$	Deterministic segment speed for a pit and load state	[km/h]
w_{starve}	Weight on the starvation term in DD scoring	[-]
w_{target}	Weight on the target-compliance term in DD scoring	[-]
w_{travel}	Weight on the travel-time term in DD scoring	[-]
w_{wait}	Weight on the queueing term in DD scoring	[-]
W_{Penalty}	Penalty weight scaling compliance deviation in the DD tuning objective	[t per percentage-point]
$\text{grade}_{\%}$	Signed road grade expressed in percent	[%]
μ_{emp}	Mean of the empirical distribution	[varies]
Δaf	Stochastic step term in the aggressiveness random walk	[-]

1

Introduction

Mining operations seek to reduce costs and improve operational efficiency continuously. In open pit mining a dominant cost contributor is the transport of material, due to the scale of fuel consumption, maintenance demand and labor costs. Therefore even modest improvements in haulage efficiency can translate into significant economic benefits.

At the same time, modern mines capture high-resolution operational data through fleet management systems and onboard telemetry. These data streams improve the operational observability and make it possible to represent the evolving system quantitatively. This quantitative representation is a prerequisite for data-driven operational control and automation.

Despite increasing data availability, the mining environment remains difficult to control because operations are inherently variable and tightly coupled through shared resources. Variability arises from inconsistent human execution of tasks as well as changing equipment and environmental conditions. Because equipment and vehicles make use of shared resources local disruptions propagate through the system easily. Improving performance therefore requires both reducing variability and increasing flexibility. In this thesis, variability refers to fluctuations in operational execution and system state, while flexibility refers to the ability to adapt operational decisions in response to those fluctuations and disturbances.

Two commonly proposed technological solutions for the variability and flexibility challenges are dynamic dispatching and autonomous vehicles. Dynamic dispatching aims to increase the flexibility by making real-time decisions on hauler allocation using the available data. Research has explored a range of methods that guide the dispatch decision making, each with its own trade-offs in interpretability, computational requirements, and robustness to stochasticity. Autonomous vehicles aim primarily to reduce the human-driven variability by executing tasks more consistently than human operators. Furthermore autonomous vehicles are able to operate with less stoppages compared to humans, as they do not require breaks and shift changes resulting in more uptime.

The Flandersbach limestone mine provides a representative setting to study these proposed solutions. Currently the operation relies on a fixed-assignment dispatch method, where haulers are assigned to a specific loader for the duration of a full shift. This limits the flexibility, as the responsiveness to disturbances during shifts is limited. The fleet at Flandersbach is fully manned, which results in performance fluctuations because of human variability. In addition, Flandersbach is considering an incremental introduction of autonomous vehicles instead of a full fleet transition. While this reduces investment risk, the implications of mixed-fleets within mining operations, especially with only a single autonomous unit, remain largely unstudied. The realized benefits of autonomy may be constrained by interactions with manual equipment and schedules, for example when breaks or shift changes pause coupled processes and limit continuous utilization.

Field evaluation of dynamic dispatching and autonomy in a live mine would provide the most realistic evidence of their operational impact. However such experimentation is risky and costly. Dynamic

dispatching can disrupt production if suboptimal allocation decisions create congestion or starve critical resources for example. Introducing autonomy can disrupt operations too and introduces safety and integration risks. For these reasons this thesis employs a data-driven Discrete Event Simulation (DES) as a cost efficient non-intrusive evaluation approach. DES captures discrete interactions, such as arrivals, resource contention, such as queues at loaders, operational constraints and the stochastic nature of processes. The DES is populated with site specific rules and constraints, and makes use of historical telemetry data from the mine. After validation the DES can serve as an offline digital twin that approximates system behavior and allows for the evaluation of the different experimental setups. Furthermore, the DES provides an environment that can be used for additional experiments too. When key assumptions change over time (for example the assumed performance parameters of autonomous equipment), scenarios can be re-evaluated efficiently by updating the relevant inputs.

Within a calibrated DES of the Flandersbach mine, the study compares a flexibility-oriented dynamic dispatching policy with the incremental introduction of a single autonomous unit in a predominantly manual fleet. Autonomous variants are evaluated by operational role (hauler, loader, or load-and-carry), focusing on production efficiency and operational stability, including queue formation and recovery after disturbances.

To keep the evaluation tractable and aligned with available data, several aspects fall outside the scope of this thesis. The DES does not model detailed autonomy workings such as perception, control, localization, or safety. Instead, autonomy is represented through performance and availability assumptions that are explicitly stated and varied across the experiments. Long-term strategic decisions are not addressed, and the results are not intended as universal performance claims beyond the studied operating rules and case constraints. Where real-world effects cannot be represented directly, simplified representations and proxies are used and are discussed as modeling assumptions and limitations.

1.1. Research Objectives and Questions

The primary objective of this thesis is to quantify and compare the operational impact of intelligent dispatching versus the introduction of a single autonomous vehicle within the specific context of the Flandersbach mine. To achieve this the following main research question has been formulated:

Main Research Question:
How do dynamic dispatching strategies and the integration of single autonomous agents compare in improving production efficiency and operational stability at the Flandersbach mine?

To guide the investigation and address the specific constraints of the operation three sub-questions have been defined:

Sub-question 1:
To what extent can a heuristic-based dynamic dispatch algorithm unlock unused capacity in the existing manual fleet while maintaining adherence to grade control targets?

Sub-question 2:
How does the introduction of a single autonomous unit (Hauler, Loader or Load & Carry) impact system performance when considering the interdependency with manual equipment?

Sub-question 3:
How do dynamic dispatch and single agent autonomy affect the balance between production performance and operational stability, as reflected by target adherence, queue formation, road interactions, and restart behavior?

1.2. Research Roadmap

This thesis is structured to move from theoretical foundations to practical application. Chapter 2 provides a literature review on dynamic dispatch algorithms and the current state of autonomous mining

equipment. Chapter 3 details the operational reality of the Flandersbach mine and establishes the constraints and rules of the case study. Chapter 4 presents the methodology, including the data-analysis pipeline, the construction of the Discrete Event Simulation, the experimental setup, and the KPI framework. Chapter 5 presents the validation and experimental results and analyzes the performance of the proposed strategies. Chapter 6 provides an interpretation of the results of the previous chapter, and discusses recommendations for further research on the topics of this thesis. Lastly Chapter 7 summarizes all the finding of the thesis, and presents the conclusion.

2

Theoretical Background

This chapter reviews previous research on material haulage in mining operations, structured around two operational objectives: reducing performance variability and increasing operational flexibility. It focuses on two intervention families discussed in the literature to support these objectives: dynamic dispatching (DD) and vehicle autonomy. First, variability and flexibility are defined in the context of mining and their operational relevance is motivated. Next, the main families of DD approaches are discussed, highlighting their characteristic assumptions, strengths, and limitations. Next, autonomy research is reviewed for three vehicle classes: haulers, loaders, and load-and-carry (LC). Finally, the chapter motivates discrete event simulation (DES) as an effective evaluation framework to compare these interventions under operational constraints, and concludes with a synthesis that motivates the methodological choices of this thesis.

2.1. Variability and flexibility in mining operations

Mining operations are frequently modeled as a cyclical truck-loader process based on average cycle times and steady operating conditions. In contrast actual execution is stochastic due to interactions between equipment, operators, and shared infrastructure, which affects throughput and utilization (Newman et al., 2010). For this thesis, it is useful to distinguish two related concepts.

First, operational variability refers to fluctuations in cycle components, such as ranging time and payload. Second, disruptions relate to unforeseen events that abruptly change the conditions of the mine, such as crusher/equipment failure. The remainder of this section focuses on both types of variability and discusses how it propagates through the system. Furthermore, this section showcases how flexibility enhancing interventions mitigate the negative effects of operational variability and disruptions (Mena et al., 2013; Moradi & Askari Nasab, 2017).

2.1.1. Variability

Operational variability stems from multiple distinct sources. A first source is heterogeneity in the fleet. Because of different capabilities and specifications, cycle times differ among vehicles of the same vehicle class (Chaowasakoo et al., 2017). A second source of operational variability is the human factor. Operators differ in their individual behavior, skill and fatigue which introduces dispersion in many important areas such as, achieved speeds, loading efficiency and ranging time (Meech & Parreira, 2011; Parreira, 2013). The third source is the state of the mining environment. Weather changes the surface conditions and, among other operational factors, causes the rolling resistance and achievable speeds to change (Ozdemir & Kumral, 2019).

Because mines are constituted of interdependent processes and the vehicles share the road network and service points (crusher, stockpiles, loaders) these fluctuations do not remain local, rather they propagate through the system (Mena et al., 2013; Newman et al., 2010). When overtaking is restricted differences in travel speed lead to truck bunching and clustered arrival at service points which leads to queue formation (Soofastaei et al., 2016). Such interaction effects are one reason why evaluation

based on deterministic averages can be optimistic (Mena et al., 2013).

Disruptions differ from operational variability because of they represent discrete abrupt changes in capacities and/or constraints, such a crusher or loader failure. Their impact can be system-wide as the unavailability of a service point causes increased usage of other service points.

Flexibility in this context concerns the ability of the system to adapt to changing conditions because of both operational variability and disruptions by reallocating resources, rerouting flows or adjusting priorities in response to the novel system state (Icarte et al., 2021; Icarte-Ahumada & Herzog, 2025). During operation flexibility is supported by physical design choices such as alternative routes, multiple dumping locations, and by responsive policies (Mena et al., 2013).

In Chapter 3, these concepts are linked to the Flandersbach operation by showcasing where variability and disruptions arise in the cycle components. Furthermore, the network, policies and destination structure providing (or limiting) flexibility are discussed, thereby motivating the dispatch and autonomy scenarios evaluated later.

2.2. Dynamic dispatching

Mine planning is typically described as a two-stage process. The upper stage sets production targets, often based on historical performance through average cycle times and expected operating conditions (Mohtasham et al., 2022). The lower stage consists of operational execution, where realized cycle times, available equipment and achieved speeds change continuously due to the operational variability and disruptions. In practice, meeting the targets efficiently requires responsive allocation decisions that adapt to the current system state (Icarte et al., 2021).

This is precisely the purpose of DD, it incorporates real-time information, and in some approaches look-ahead, to allocate haulers to loading and dumping locations rather than relying on fixed assignments (Noriega et al., 2025). Due to advancements in sensor technologies and the increased amount of data stored the quantity and resolution of real-time information has increased (Ozdemir & Kumral, 2019). This enables DD methods to characterize the system state in more detail and incorporate multiple objectives, but also increases the complexity of the DD policies (Firoozi et al., 2024; Noriega et al., 2025).

The remainder of this section reviews five families of DD approaches, and discusses their core mechanism, the information they require, their objectives and reported outcomes and their key limitations. Table 2.1 summarizes these findings and highlights the trade-offs per DD family.

2.2.1. Heuristics

Heuristic DD forms the earliest and most widely deployed family of dispatch policies (A. Y. Voronov et al., 2023). The allocation decision is made by ranking destinations based on simple decision criteria and selecting the highest ranking option. Many implementations use a 1-hauler-for-N-loaders model. In this setup the dispatcher calculates the costs and benefits for every possible loader and sends the hauler to the loader with the highest rank, once a hauler becomes available for its next assignment (Smith et al., 2021). This procedure is computationally inexpensive and operationally transparent, as the decision is reduced to evaluating a simple score across a limited set of loaders and the inputs that compute this score are explicit (Alarie & Gamache, 2002).

Heuristic rules rely on immediately available state variables and target tracking measures. Examples include queue lengths and ratios that reflect deviations from predefined targets (Noriega et al., 2025). The expected travel and service times used for ranking are computed via deterministic values, often based on historical averages. Therefore heuristics do not represent the variability of cycle components or condition-dependent changes (Alarie & Gamache, 2002; Smith et al., 2021). Because of the simplicity of the state and system representation this approach allows for frequent re-evaluation, in contrast to forward simulating approaches or global optimization strategies (Alarie & Gamache, 2002).

Quantitative studies report that heuristic DD yields better production KPIs compared to fixed assignment methods. In a 6.5 hour Sungun copper mine simulation with 19 haulers and 4 loaders, a multi-stage heuristic hauled 22,800 t compared to 19,312 t in a fixed assignment setup and 22,600 t when making use of DISPATCH® (Mohtasham et al., 2022). DISPATCH® is a commercial dispatch system that

implements real-time allocation rules and is commonly used as a reference baseline in mine dispatch studies (Alarie & Gamache, 2002; White & Olson, 1986). However, the OpenMines research shows that overly simple heuristics perform poorly compared to fixed assignment (Meng et al., 2024). In this simulation the shortest-queue policy produced 13,332 t and the fixed assignment policy produced 14,910 t.

The OpenMines simulation demonstrates a key limitation of heuristic methods, myopia. The ranking rules optimize the assignment without anticipating how this assignment influences the future system state, which can result in non-optimal decisions because of the shared resources (Alarie & Gamache, 2002; Smith et al., 2021). Relatedly, because the ranking is based on historical deterministic values rather than distributions, the policy does not incorporate the inherent variability in its ranking. Furthermore, heuristic performance is context-dependent because dispatching criteria implicitly assume a dominant bottleneck. In undertrucked conditions, for example, the primary objective is often to avoid shovel idleness, whereas in over-trucked conditions queues dominate, rendering the gains from such a heuristic ineffective (Alarie & Gamache, 2002; Chaowasakoo et al., 2017). Lastly, mining involves competing objectives such as total tonnage and target rates per location resulting in a multi objective problem. However, single criterion heuristic are not suited to incorporate more than one objective (Mirzaei-Nasirabad et al., 2023).

To partially address these limitations, more modern approaches have moved from single-attribute to multi-attribute heuristics (Munirathinam & Yingling, 1994). By incorporating multiple factors into a weighted sum, competing objectives can be balanced within the ranking. Additionally some strategies incorporate small look-aheads, to counter the negative effects of the myopic nature of the heuristic methods (Subtil et al., 2011). However, both these solutions remain partial. The introduction of weighted sums in the ranking introduces trade-offs regarding the competing objectives. To preserve the computational efficiency of heuristics the look-aheads need to be small and thus do not capture the full impact the decision has on the future system state. Therefore, both solutions do improve on the simple heuristics but suffer from the same shortcomings.

Overall, heuristic DD provides fast and interpretable decisions, but its myopic structure can yield suboptimal outcomes. Moreover, because these dispatch rankings are typically based on deterministic values rather than stochastic representations, heuristics neglect the impact of variability (Alarie & Gamache, 2002; Smith et al., 2021).

2.2.2. Mathematical programming

The family of mathematical programming (MP) techniques formulate dispatching as a system of equations and inequalities. Within this system decision variables, constraints and objective function(s) represent the real operation logic such as, capacity limits and production requirements (Smith et al., 2021). Instead of ranking possible destinations, MP selects assignments by solving the system of equations model. Solving the system of equations is a computationally expensive, and thus time consuming, procedure. Therefore MP is naturally suited for upper-stage planning, where time is not constrained and flow rates and production requirements are the objectives. When time is not constrained MP can yield mathematically optimal solutions, however in the lower-stage real time decisions are required (Smith et al., 2021). Literature demonstrates that MP can also be used for the lower-stage. In this configuration, the model must be adapted to meet these latency requirements by limiting model size, horizon length, or optimality (Mirzaei-Nasirabad et al., 2023; Smith et al., 2021).

To achieve this reduced model size and complexity lower stage MP dispatching typically makes use of simplified formulas and/or a reduced scope. One approach uses mixed-integer programming over a short time frame, producing allocation decisions while periodically re-optimizing as the system state evolves (Smith et al., 2021). Another approach reduces the scope to a 1-hauler-for-N-loaders model. Each time a hauler needs a new destination the model optimizes only that immediate assignment via binary decision variables, thus the consequences for the remainder of fleet are not evaluated (Mirzaei-Nasirabad et al., 2023).

In the same Sungun copper mine, a multi-objective MP dispatch method increased production from 19,313 t under fixed allocation, to 21,880 t (Mirzaei-Nasirabad et al., 2023). Beyond tonnage, MP-based dispatching is also evaluated on target compliance objectives. Smith et al. show that using a

bounded-time MIP solution method can significantly reduce ore-quality target violations and increase the fraction of time ore quality stays within a 1% band of the target under strict solve-time limits (Smith et al., 2021). For a 20-minute horizon, the percentage of time ore quality remained within a 1% band of the target increases from 70.0%, for a reference solution that does not optimize ore quality, to 95.0% (Smith et al., 2021).

The dominant limitation of MP for DD is computational efficiency and associated latency. Exact fleet-wide formulations scale poorly with the problem scope, creating significant solve time increases (Smith et al., 2021). Consequently, MP approaches are myopic by design as the time horizon considered and amount of considered alternatives are reduced, resulting in feasible decisions rather than globally optimal decisions. Secondly, uncertainty is no incorporated in this approach. Instead, it uses deterministic values or historical averages for cycle components, this renders computed solutions suboptimal as conditions deviate from the assumptions (Noriega et al., 2025). A further limitation is objective specification. In practice dispatching is multi-objective, but standard MP formulations require collapsing competing goals into a single scalar objective, which forces explicit trade-offs that may not reflect operational preferences (Mirzaei-Nasirabad et al., 2023; Temeng et al., 1997).

To counteract these limitations several variants have been designed. Goal programming enables the treatment of multiple competing objectives by minimizing deviations from predefined targets rather than optimizing a single objective (Moradi Afrapoli et al., 2017; Noriega, 2023; Temeng et al., 1997). Another extension separates the dispatch decision into two stages by distinguishing loaded-truck dispatching to destinations and empty-truck dispatching to loaders (Wang et al., 2023). This decomposition allows the second assignment to be updated after the first leg is completed, improving responsiveness to changes in system state (Wang et al., 2023).

Overall, MP provides a structured way to encode dispatching constraints and objectives as a system of equations. It can be applied to lower-stage DD when adapted to meet latency requirements through restricted scope. However, these reductions and deterministic input assumptions limit the extent to which MP captures stochasticity and interaction effects in real operations. This motivates evaluating MP-based DD interventions in interaction-aware frameworks, such as DES, when comparing DD families.

2.2.3. Metaheuristics

Metaheuristics are search procedures that generate near-optimal solutions within practical time frames, by exploring the solution space in an efficient way focusing on promising regions (Moradi & Askari Nasab, 2017). Common approaches include genetic algorithms, tabu search, ant colony optimization and particle swarm optimization (Iqbal et al., 2021; Noriega & Pourrahimian, 2022). Although the precise workings of these methods differ, they rely on the same principle. Only a select set of candidate solutions are evaluated and iteratively improved (Aoshima, 2025). A central trade-off is balancing the exploitation and exploration phase. The exploitation part is where the search algorithm focuses on a promising region within the solution space, and the exploration part is used to evaluate new parts of this solution space (Peres & Castelli, 2021).

In DD, metaheuristics evaluate candidate decisions or candidate policy parameters in regard to operational performance. This is typically done through simulation-based optimization (SBO), where a calibrated simulation model represents the mine in detail. By representing the fleet, service points and interactions within the mine candidates are scored under congestion and queuing effects (Abolghasemian et al., 2020; Mena et al., 2013; Ozdemir & Kumral, 2019). Because repeated simulation is computationally expensive the SBO is typically used as offline or periodic optimizer that tunes a policy. After tuning the execution relies on a lightweight computation, instead of a full simulation at every dispatch request (Abolghasemian et al., 2020).

Metaheuristic-based SBO can represent multiple objectives through weighted cost functions or target-deviation terms (Mena et al., 2013; Ozdemir & Kumral, 2019; Soleymani Shishvan & Benndorf, 2018). Quantitative studies report significant production improvements compared to baseline dispatching strategies in simulations. Ozdemir and Kumral report that their simulation-based optimization framework increases total production by 6.0 kton per shift, corresponding to a +9.4% improvement over the baseline dispatching system (Ozdemir & Kumral, 2019). Abolghasemian et al. report a production increase

of 10,000 t (+21%) relative to the baseline strategy using a two-phase SBO approach supported by metaheuristics (Abolghasemian et al., 2020). These findings illustrate the usefulness of SBO and metaheuristics for DD. Metaheuristics efficiently search the policy parameters solution space, while the simulation captures the interaction effects within the system that are difficult to represent in deterministic models.

Metaheuristics do not guarantee global optimality and can converge to local optima if the exploration and exploitation is in imbalance (Peres & Castelli, 2021). Performance highly depends on this balance, the defined objective and parameter settings (Moradi & Askari Nasab, 2017). Another limitation is the computational burden of this technique. Repeated simulation is time-consuming, motivating offline tuning or periodic re-optimization instead of simulating per dispatch request (Abolghasemian et al., 2020). Finally model fidelity is of high importance, if the simulator does not reflect real operational variability and constraints, tuned policies may not be transferable to the real situation (Saleem, 2025).

To address the runtime limitations noted above, practical SBO frameworks separate policy tuning from policy execution. Simulation is used offline or uses update intervals to tune allocations or parameterized dispatch rules, after which the real-time layer applies fast scoring rules or lightweight optimization to generate decisions (Abolghasemian et al., 2020; Ozdemir & Kumral, 2019).

Overall, metaheuristics combined with SBO provide a flexible way to search for high-performing DD policies, and they can incorporate stochasticity through simulation during the tuning phase. However, their effectiveness hinges on simulation fidelity, motivating careful evaluation of the tuned policies.

2.2.4. Multi Agent Systems

In contrast to the other families described in this chapter Multi Agent Systems (MAS) implement DD through decentralized decision-making (Icarte et al., 2020, 2021). In the mining setting, the agents within this model represent haulers, shovels/loaders, crushers and stockpiles (Icarte et al., 2020). Dispatching is executed via a Contract Net Protocol (CNP). This protocol operationalizes DD as a negotiation (bidding) process (Icarte et al., 2020, 2021). In this process a shovel agent issues a call for proposals, hauler agents submit proposals, and the shovel selects a proposal based on expected timing and operational criteria (Icarte et al., 2020, 2021). Other techniques only issue a next-destination decision, whereas MAS constructs a truck schedule (a sequence of operations with timing and locations), enabling coordinated execution over a pre-defined period (Icarte et al., 2020).

Under the CNP agents act independently using local data during the call for proposals and bidding process (Icarte et al., 2020, 2021). Proposals include expected arrival and operation times derived from truck parameters such as capacity, empty/loaded velocities, spotting, and unloading times (Icarte et al., 2021). These estimates are treated deterministically via fixed parameter values, and routing is simplified through shortest-path traversal rather than explicitly modeling micro-traffic behavior (Icarte et al., 2021).

MAS studies report outcomes on both production and cost-related KPIs. In a Chilean open-pit case study, the baseline results are taken from actual shift data from the mine using DISPATCH (Icarte et al., 2021). MAS achieves small positive production differences, with transported material improving by approximately +0.20% to +0.63% (Icarte et al., 2021), compared to the baseline. The larger reported benefit is hauler cost reduction. MAS reduces this cost by decreasing hauler operating hours by approximately 15% to 18% while maintaining or slightly increasing transported tonnage. (Icarte et al., 2021). MAS also reports favorable computational performance relative to exact optimization. In a setup with 40 trucks and 5 shovels, MAS generates a shift schedule in 2.45 minutes, whereas an exact mathematical programming formulation solved with CPLEX fails to find a solution within 180 minutes (Icarte et al., 2021). When a disruption breaks the current schedule, MAS is able to respond in two distinct ways. For local events, it performs a schedule repair that cancels only the affected assignments and fills the resulting gaps. More severe cases trigger a complete reschedule in which agents cancel all their assignments and renegotiate a full new schedule (Icarte et al., 2021). In the reported large-fleet breakdown experiment, schedule repair required 15 seconds, whereas complete rescheduling required 785.5 seconds (13.1 minutes), indicating that both responses are computationally feasible within operational time scales (Icarte et al., 2021).

A limitation of MAS is the near-optimal solutions found with the CNP (Icarte et al., 2021). The approach

also relies on a simplified operational representation with deterministic estimates, which limits its capabilities to handle stochastic variability and traffic effects (Icarte et al., 2021). In addition, MAS is reactive as it repairs or reschedules when conflicts or failures occur (Icarte et al., 2021).

Later works on MAS integrate reinforcement learning into its scheduling framework. This addition improved schedules compared to the non-learning baseline in simulated scenarios (Icarte-Ahumada & Herzog, 2025).

Overall, MAS offers a responsive DD strategy that can reduce operating costs while maintaining production. The MAS strategy is able to generate schedules within practical run times too. However, its reliance on simplified timing and routing assumptions and its reactive repair logic motivate evaluation in interaction-aware frameworks that explicitly represent congestion, variability, and disruption propagation.

2.2.5. Reinforcement Learning

Reinforcement learning (RL) comprises a family of techniques in which an agent learns an efficient dispatch policy through trial and error (Noriega et al., 2025). The objective is to maximize cumulative reward over time rather than the immediate reward, which encourages decisions that account for downstream consequences (Noriega et al., 2025). This trial and error process is typically performed in simulation because RL requires many interactions, additionally real-world trials are costly and can be unsafe (Eriksson et al., 2024). The learning process is formalized via a Markov decision process, a mathematical framework to model sequential decision making (Barde et al., 2019). The most important assumption of this framework is the Markov property, which is defined as: the current state of the environment contains all the information needed to make an optimal decision (Noriega et al., 2025). For each decision the agent observes the current state, takes a dispatch action, receives a reward, and transitions to a new state (Barde et al., 2019; Noriega et al., 2025). After training, dispatch decisions are produced in real time, because evaluating the learned policy is computationally fast (de Carvalho & Dimitrakopoulos, 2021; Noriega et al., 2025).

RL-based DD approaches represent this mine state as a vector capturing operational conditions relevant for the dispatch decision such as, queue sizes, vehicle positions and plant feed requirements (de Carvalho & Dimitrakopoulos, 2021; Noriega et al., 2025). Actions correspond to the assignment of a hauler to a destination, while rewards are constructed to balance throughput and target adherence against penalties for waiting and deviations from the requirements (de Carvalho & Dimitrakopoulos, 2021; Noriega et al., 2025).

De Carvalho and Dimitrakopoulos compare a deep RL dispatch strategy based on Double Deep Q-learning (DDQN) against a greedy shortest queue and a fixed policy baseline. Over a five day horizon their DDQN policy increases copper recovery by 16% compared to the fixed policy and 12% to the shortest queue policy. For gold, the DDQN policy increases recovery by 20% relative to the fixed policy and by 23% relative to the greedy policy (de Carvalho & Dimitrakopoulos, 2021). Noriega et al. show that a DDQN-based dispatching agent meets loader and shift targets in their simulation and is generalizable to smaller fleets, maintaining targets with 26 trucks after being trained on a 34-truck fleet (Noriega et al., 2025).

RL performance depends on the fidelity of the training environment. If the simulation does not reflect how the mine responds to dispatch decisions, the learned policies can be biased or fail to be transferrable to reality (Noriega et al., 2025). A further important limitation is interpretability. RL policies are black boxes, which complicates validation and adoption in safety-critical environments such as mines (Meng et al., 2024). Finally, the training phase is computationally intensive because of the extreme amounts of iterations the trial-and-error phase requires (Eriksson et al., 2024; Noriega et al., 2025).

For DD, deep RL is often implemented with Deep Q-Network (DQN) methods, where a neural network approximates the state–action value function to handle large state spaces (Noriega et al., 2025). Standard DQN can lead to systematic overestimation, which can destabilize learning and reduce policy quality (Noriega et al., 2025). DDQN mitigates this limitation by decoupling action selection from action evaluation during the target computation, reducing overestimation bias (Noriega et al., 2025). Practical implementations typically also use experience replay and a target network to stabilize training (Noriega et al., 2025).

Overall, RL enables adaptive DD policies that can anticipate downstream effects through reward-driven learning and has demonstrated substantial gains over myopic baselines. However, it relies on high-fidelity simulation, heavy training requirements. Limited interpretability motivates careful evaluation under realistic variability, congestion, and disruptions when comparing DD families.

2.2.6. Comparison

Every DD family reviewed above has its strengths and weaknesses. In Table 2.1 an overview, which summarizes these strengths and weaknesses and highlights performance outcomes reported in literature, is listed. Based on the analysis presented in the previous section, the three main trade-offs regarding the selection of a dispatch strategy are:

- **Optimization scope vs. decision latency:** Methods that optimize over larger horizons or more coupled decisions can improve system-level performance, but they typically increase computation time for training or for producing a dispatch decision. Consequently, global scope is often traded for real-time responsiveness through reduced horizons, reduced decision scope, or offline tuning.
- **Performance vs. interpretability:** High-performing adaptive methods (notably deep RL) can learn complex dispatch behavior, but often provide limited transparency into why specific decisions are made. This complicates validation and adoption in safety-critical settings.
- **Deterministic vs. stochastic representation:** Simpler deterministic models are easier to construct and solve, but they can neglect operational variability and interaction effects. Methods that evaluate policies in simulation can incorporate stochasticity and congestion, but their results depend on model fidelity and incur higher computational cost.

2.3. Autonomy

A current shift in mining operations is the implementation of autonomous vehicles (Kaur, 2024). These vehicles operate without a human driver, which can yield operational benefits while also introducing technical and organizational risks. The benefits and drawbacks all stem from the complete removal of the human element of operating a vehicle (Moniri-Morad et al., 2024). The Global Mining Guidelines Group emphasizes a holistic approach is needed that integrates people, processes and technology, for the successful implementation of autonomous vehicles (Global Mining Guidelines Group, 2024). In this thesis, this broader perspective is acknowledged. However, the scope is limited to technical and operational aspects that can be represented in a simulation model and their implications for performance. The following paragraphs will first outline the general benefits and drawbacks of autonomous vehicles in the mining industry. It then discusses three autonomous vehicle classes that are either deployed or in pilot phases within the mining industry.

2.3.1. General Benefits and Drawbacks of Autonomous Vehicles

Removing human operators from vehicles within the mine can improve safety by reducing worker exposure to high-risk tasks and environments (Dadhich, 2018). Mining environments are hazardous due to the presence of heavy vehicles, limited visibility, and geotechnical threats. Collisions involving large mobile equipment can cause severe injuries or fatalities, workers may be exposed to dust and noise and falling rocks. Accordingly, reducing human presence in these areas is reported as a pathway to fewer injuries and fatalities (Y. Voronov et al., 2020). Humans can be the source of unsafe situations too. Humans, in contrast to robots, suffer from fatigue, distraction and inattention or may not respect the safety rules (Moniri-Morad et al., 2024). Automation can mitigate these safety risks, however it may also introduce new failure modes that require appropriate controls.

The second driver of autonomous vehicle implementation is increased productivity. Case studies report productivity increases ranging from 15 to 20% for autonomous setups (David Bird, 2019). Robots do not need breaks and shift changes. For a mine that operates continuously the removal of the breaks and shift changes can amount to 700 hours more operating time per year per hauler (Long et al., 2024). Another way production can be increased with autonomous vehicles is their increased consistency and predictability compared to human operators (Corke et al., 2008). Because of distractions, fatigue and skill differences human operators perform tasks with greater variance compared to robots. Variance reduction leads to decreased fuel usage, less frequent maintenance of the vehicles and a more pre-

Table 2.1: Comparison of Dynamic Dispatch (DD) strategy families. Note: execution speed refers to the time required to generate a decision when a truck requests a destination.

Strategy Family	Core Mechanism	Execution Speed	Variance Handling	Transparency	Main Limitation	Reported Yield
Heuristics	Ranking and rule-based assignment	Instantaneous	Low (deterministic values)	High	Myopic decisions and limited multi-objective control	22,800 t vs. 19,312 t fixed and 22,600 t vs. DISPATCH (Mohtasham et al., 2022)
MP	Constrained optimization with reduced scope for real-time use	Fast if reduced-scope	Low (often historical averages)	High	Latency and scalability for larger scopes	21,880 t vs. 19,313 t fixed (Mirzaei-Nasirabad et al., 2023)
Metaheuristics / SBO	Offline or periodic tuning with simulation-based evaluation	Slow (tuning) and fast (execution)	High during tuning (simulation captures interactions)	Moderate	High computational burden and model-fidelity dependence	+6.0 kton/shift (+9.4%) (Ozdemir & Kumral, 2019) +10,000 t (+21%) (Abolghasemian et al., 2020)
MAS	Decentralized negotiation via CNP producing schedules	Moderate	Moderate (reactive repair with simplified routing)	Moderate	Near-optimal solutions and simplified timing and routing assumptions	Hauler cost –15% to –18% with +0.20% to +0.63% transported material vs. DISPATCH baseline (Icarte et al., 2021)
RL	Learned policy trained in simulation (e.g., DDQN)	Instantaneous	High if trained in stochastic simulation	Low	Training intensity and limited interpretability	+16% Cu and +20% Au vs. fixed policy +12% Cu and +23% Au vs. greedy policy (de Carvalho & Dimitrakopoulos, 2021)

dictable operation (David Bird, 2019). With smoother and more consistent operation, autonomy can reduce fuel consumption and slow wear of critical components such as tires and gearboxes (Gustafson, 2011). The last way an autonomous vehicles can optimize mining processes is because it allows for different mine designs. Because operator exposure constraints are reduced in a fully autonomous setup, certain geometric and operational design constraints such as road widths or safety distances can potentially be optimized, When the vehicles know the position of other vehicles it is possible to make narrow corners even under limited line-of-sight conditions too. This allows for enhanced resource extraction, because former constraints in the mine plan have been reduced/removed (Rogers et al., 2019).

The main drawbacks of autonomous vehicles lie in significant upfront investments (Wang et al., 2023) (Saleem, 2025). A communication network must be setup, and new trucks must be purchased or old trucks have to be retrofitted with software to be able to drive autonomously. In unexpected cases robots may be unable to operate as they are not as adaptable as humans. Robots do not posses the same level of reasoning, judgment and foresight as humans (Rogers et al., 2019). In rare situations (unexpected obstacles or atypical ground conditions), autonomous systems may behave conservatively and require human intervention, reflecting limits in adaptability compared with human judgment (Rogers et al., 2019). The workforce needs to undergo significant changes too (Firoozi et al., 2024) (David Bird, 2019). This transition is commonly associated with changes in job roles, including reduced demand for routine driving tasks and increased demand for technical and supervisory roles. These new roles require far more digital skills and understanding of autonomous vehicles in general (Firoozi et al., 2024).

2.3.2. Autonomous Haulers

Autonomous haulage systems (AHS) are the most mature class of autonomous mobile equipment in mining, with commercial deployment in surface operations. However, global adoption of AHS remains limited. In Figure 2.1 the amount of autonomous haulers per country is displayed. The figure indicates that adoption is still concentrated in a limited set of countries, and the adoption increases over the reported period (MiningTechnology, 2022). In Europe a limited number of autonomous haulers have been employed in Ukraine, Sweden and Norway, suggesting local context may influence adoption in the region (MiningTechnology, 2022). The maturity of AHS compared to other mining vehicle types is often attributed to the comparatively simple haulage task, whereas loading and LC involve complex interactions with fragmented material (Tampier et al., 2021).

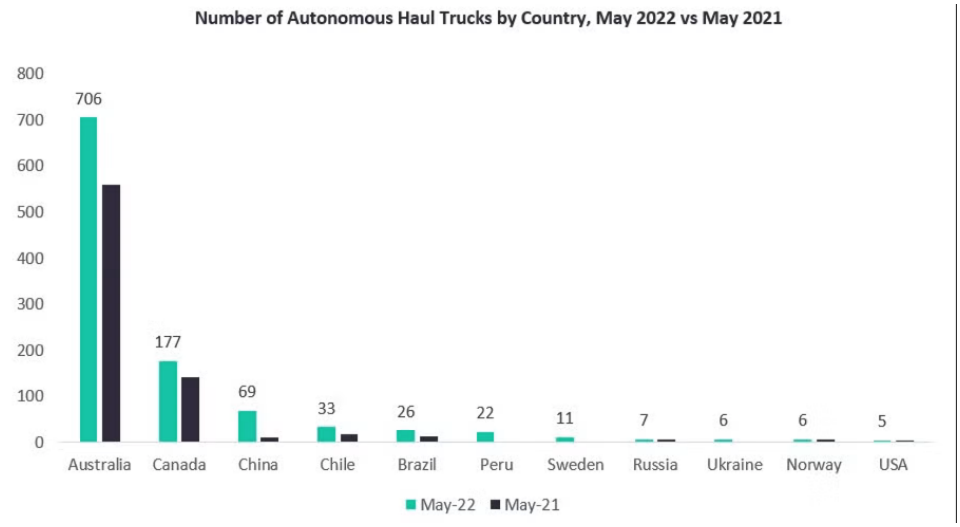

Figure 2.1: Adoption of autonomous haulers per country in 2021 and 2022. (Source: (MiningTechnology, 2022))

Implementing AHS requires several technical prerequisites to be met. AHS relies on high-accuracy localization methods on the order of decimeters (Parreira, 2013). Localization alone is insufficient for detecting obstacles or employees present in the mine. Therefore AHS systems are fitted with Lidar or Radar sensors to detect these objects/employees (Corke et al., 2008) (David Bird, 2019). The fleet is coordinated by wireless comunnication (Wi-Fi, 4G, 5G) (Gaber et al., 2021) (Saleem, 2025). This network provides the information for centralized coordination, including hauler dispatching and

destination assignments. A commonly noted vulnerability stems from this centralized coordination and wireless communication. Failures in one of these systems leads to operation wide interruptions (Y. Voronov et al., 2020).

Although AHS typically increases safety it can impose new hazards too (Moniri-Morad et al., 2024). A substantial portion of these hazards stem from the difficult human-robot interaction (Burgess-Limerick et al., 2024) (Parreira, 2013). Therefore a key consideration for AHS implementation is the frequency of these human-robot interactions. There are three approaches to this implementation: a complete replacement, a staged replacement where the autonomous haulers operate on isolated routes, and a mixed system where human operators and robots operate alongside each other (Parreira, 2013). Recent reviews note that mixed operations are discussed as a safety concern, but are still rarely analyzed explicitly as a distinct hazard category in the literature (Moniri-Morad et al., 2024). Interaction frequency is lowest under full segregation, increases under staged/isolated deployment, and peaks in mixed human–robot operations. The risks of the human-robot interaction are caused by both the robots and the humans (Parreira, 2013). The robots have a hard time dealing with human behavior, as robots exhibit more rigid behavior relative to human drivers, which can create mismatches in expectations during non-routine situations. On the other hand, human operators may misjudge the capabilities and limitations of autonomous vehicles, especially regarding conservative decision-making and limited adaptability (Long et al., 2024). Importantly the interactions between vehicles within mines are not homogeneously distributed throughout the mine, but are concentrated at operational hotspots such as intersections, loading areas and dumping zones (Moniri-Morad et al., 2024). For a conventional fleet these places are the most dangerous too. Md-Nor et al report that 40 of 111 hauler related fatalities (≈36%) occurred at dump sites and loading areas, while the amount of time hauler operators spend at these zones is much lower (Md-Nor et al., 2008). This disproportionate number of fatalities in these low speed zones stems from the tight layout and vast amount of interactions taking place. This pattern is consistent with recent modeling work showing that high-risk areas cluster around high-traffic intersections and constrained interaction zones (Goli et al., 2025). In mixed operations, these hotspots may become more critical due to human–robot interaction challenges.

The two major AHS suppliers are Komatsu and Caterpillar(Corke et al., 2008). Precise data on variance reduction for autonomous haulers is not publicly available, due to its commercial sensitivity(David Bird, 2019). Nevertheless several KPIs of AHS are reported, these values are often attributed to reduced variability and continuous operation. Reported effective utilization is around 60% for manned fleets, increasing to approximately 75–90% under autonomous operations (A. Y. Voronov et al., 2023) (David Bird, 2019). In a case study of Rio Tinto the autonomous fleet operated 700 hours more compared to a manned one (Long et al., 2024). In another case study AHS achieved a 15% reduction in the load and haul unit costs, a result from around 6% lower fuel consumption and reduced component wear of around 7% (A. Y. Voronov et al., 2023).

2.3.3. Autonomous Loaders and LC

Autonomous loaders and LC are less mature than autonomous haulage systems (AHS) and are predominantly reported at teleoperation or semi-autonomous levels. Currently, both these vehicle types have not been deployed commercially, but the first pilots are being conducted (Du et al., 2025). Both Liebherr and Volvo currently conduct tests with their autonomous wheel loader prototypes (Equipment Journal, 2021; Liebherr-Werk Bischofshofen GmbH, 2025). Therefore data on these topics is more scarce compared to AHS. A larger share of research has been conducted for load haul dump machines (LHD), which operate underground, but suffer from similar challenges as loaders and LC (Fernando & Marshall, 2020). Findings from this research help clarify the main reasons why this form of autonomy is less mature than AHS.

A key difference between loaders, LC and LHD compared to haulers is that these vehicles execute the loading part of the haulage cycle (Dadhich, 2018). Digging and bucket filling is widely identified as the principal bottleneck regarding autonomy for the loading stage (Aoshima, 2025; Tampier et al., 2021). The unpredictable forces, complex geometry and fragmentation differences between the rocks make the filling of the bucket complex for autonomous vehicles (Dadhich, 2018; Tampier et al., 2021). Human operators combine their vision, hearing and their feeling of vibrations to adapt their dig strategy to the characteristics of the rock pile (Dadhich et al., 2019). Autonomous loaders, LC and LHD have

to rely on cameras, Lidar/Radar systems and internal measurements (Zhang et al., 2021). These systems primarily observe the surface and do not reveal subsurface properties that strongly affect the digging process (Fernando & Marshall, 2020). As a result, this loading stage remains challenging, and loading performance of experienced operators has not yet been matched by autonomous vehicles (Backman et al., 2021). Shared-autonomy concepts have been tested too. In this setup the digging is done autonomously, but when the dig fails a human operator takes over (Tampier et al., 2021). In field studies of automated LHD operation, operating-environment disturbances such as oversized boulders account for a large share of stop occasions causing idle time, reported at approximately 75% of the total idle time in one study (Gustafson, 2011).

The research focus on LHD autonomy is partially explained by underground operational constraints and safety considerations, as LHDs are primarily used in underground operations. Underground mines operate under more stringent safety regulations compared to open pit mines particularly during/after blasting. After blasting, conventional loading cannot resume until the blast area has been ventilated thoroughly to clear out toxic gases and restore safe working conditions (Gustafson, 2011). Automated LHD operation can resume earlier, reducing this post-blast downtime (Gustafson, 2011). The research on autonomous LHD indicates that repeatedly following an identical path can accelerate localized road degradation, motivating path variation in automated operation (Gustafson, 2011).

Because autonomous loaders and LC are at prototype or pilot stages, direct empirical evidence on their interaction risks in mixed-fleet mining operations is non existent. However loaders and LC operate in zones with dense interactions, particularly during the short-cycle loading and dumping at crusher areas of the LC (Dadhich et al., 2016). LC must navigate these high-interaction zones while maneuvering, digging and dumping, which increases exposure to traffic conflicts and constrained-geometry interactions. Therefore, autonomy concepts for loaders and LC must be evaluated with explicit attention to interaction density and constrained operating geometry.

2.4. Discrete Event Simulation and Digital Twins

To effectively analyze complex mining systems, selecting an appropriate modeling technique is critical. Discrete Event Simulation (DES) is defined as a modeling approach where a system changes state at discrete points in time, such as a truck arrival or a shovel load, allowing the simulation to skip the intervals between these events (Bernardi et al., 2020; Chaowasakoo et al., 2017). Due to its maturity and effectiveness, DES has been classified as a “Star” technique in mining literature(Noriega & Pourrahimian, 2022). It is extensively adopted because it represents the static structure of a system (network and capacities), while capturing dynamic and stochastic operational behavior over time (Bernardi et al., 2020; Ozdemir & Kumral, 2019).

Furthermore, a calibrated DES can be used not only as a generic model, but also as the simulation engine for an online Digital Twin (Noriega & Pourrahimian, 2022). An online Digital Twin is a continuously updated virtual representation of a physical system used to mirror behavior for prediction and optimization (Noriega & Pourrahimian, 2022). By populating a DES model with large historical data and real-world parameters, it is possible to simulate equipment behavior and interactions, creating a virtual counterpart of the physical mine suitable for evaluating different operational changes (Noriega & Pourrahimian, 2022; Saleem, 2025).

Deterministic models, such as MP, can overestimate production capabilities when they ignore operational variability and interaction effects (Bakhtavar & Mahmoudi, 2020). Mining operations are inherently stochastic: cycle times, breakdowns and payloads vary significantly over time and under changing operational conditions. DES is essential for this research because it captures these stochastic behaviors and the complex interactions of traffic, such as truck bunch effects (Noriega et al., 2025). These interactions impact the performance of dynamic dispatch and autonomous haulage systems significantly. Furthermore, testing operational changes such as autonomous load-and-carry in a live operation is costly and risky. A high-fidelity DES allows for the evaluation of these new fleet configurations to quantify benefits before physical implementation.

Finally, once calibrated, a DES model is highly flexible. Assumptions and parameters can be updated as improved data become available, and alternative operating procedures can be evaluated with new scenario runs. This makes DES particularly useful in settings where system knowledge evolves over

time and where proposed changes must be assessed before implementation.

2.5. Synthesis and Strategy Selection

This literature review indicates that DD can be implemented through multiple methodological families, each with clear trade-offs between solution quality, interpretability, computational requirements, and robustness to variability. RL renders complex decision policies, but offers limited interpretability and depends on extensive training and careful generalization to unmodeled situations. MP can provide optimal or near-optimal solutions under simplified assumptions, but is often computationally demanding for real-time execution and neglects stochastic variability and system interactions. In contrast, heuristic dispatch rules are computationally efficient and interpretable, yet can be myopic when they rely on limited state information. Metaheuristics and SBO can improve heuristic performance by tuning parameters offline, but this introduces additional computational burden during the design phase. MAS implement DD via decentralized negotiation, and can construct short-horizon schedules. However, implementations often rely on deterministic timing estimates and simplified routing and thus do not capture traffic interactions and stochasticity.

A central gap emerging from the reviewed dispatch literature is that dispatch policies are frequently assessed in environments that simplify or neglect interactions such as congestion and bunching despite evidence that implicate these interactions affect production significantly. This motivates a system-level evaluation approach in which interaction effects are represented. Therefore, this thesis adopts a calibrated DES as the evaluation environment and selects an offline SBO approach to tune a weighted, heuristic DD rule. Offline tuning via a metaheuristic retains fast real-time execution while repeated simulations filter out myopic policies. The resulting DD strategy remains interpretable and makes the trade-offs between competing objectives explicit (production and compliance with operational targets).

Regarding autonomy, AHS represent the most mature and only commercially deployed form of mobile equipment autonomy in surface mining, yet global adoption remains limited. Deployments are concentrated in a small number of countries, suggesting that local operating conditions and economic context influence adoption. For Flandersbach, an important contextual factor is that autonomy benefits tied to labor substitution may differ from those reported in highly remote or harsh environments. In addition, autonomous loaders and LC remain at prototype or pilot stages, and the literature contains limited quantitative evidence on their system-level impacts. Consequently, system-level studies that evaluate the introduction of autonomous loaders or LC have not yet been conducted. AHS has been investigated more, but system-level studies remain very limited.

A second central gap in the autonomy literature is the limited quantitative understanding of incremental deployment in mixed fleets, where autonomous and human-operated equipment coexist. This gap is particularly relevant because mines often pursue staged adoption rather than full-fleet replacement. Therefore, this thesis evaluates autonomy in the incremental form most relevant to the case study: introducing a single autonomous unit into an otherwise conventional fleet. This design also reflects the practical constraint that single-unit autonomy may not deliver the same benefits as full-fleet autonomy. By comparing DD policies against incremental autonomy scenarios within the same calibrated DES, the thesis directly addresses the gap at the intersection of dispatch and autonomy.

3

Flandersbach

This chapter describes the current operational of the Flandersbach mine. First the physical layout of the mine is presented. Subsequently two concise tables are given which list the vehicles present at the mine and displays their important specifications. The important processes and rules are then discussed. Lastly, a timeline is presented of a typical day. Several schematic drawings are also presented that support the interpretation of the system layout and. The information in this chapter provides the basis for the input and mechanics of the digital model, outlined in 4

3.1. Physical Layout of the Mine

This section describes the infrastructure of the mine. In Figure 3.1 an aerial map of the Flandersbach mine is depicted. In this map the locations of the crushers, breakroom, gas station and the tunnel is depicted. The different benches can be seen from the figure too. The location of the benches changes over time as blasting, excavating and hauling progress. The crushers are connected to the plant via a single conveyor belt. Consequently, conveyor belt failure causes an inability to dump at both crushers which impacts the production significantly. The tunnel forms a physical bottleneck. It is too narrow for two trucks to traverse in opposite direction, and is the only single-lane road in the mine. Figure 3.1 shows the mine comprises two pits. These pits are both operational but the characteristics of these pits differ. The Rohdenhaus pit has a higher ore grade than Silberberg, which has significant implications on the workings of the mine. The Rohdenhaus pit is the deepest pit, with a minimal Z coordinate of -50. The Silberberg pit is the furthest away, but is less deep with a minimal Z coordinate of 60. The mean grade of the roads in the road network is comparable between the pits and amounts to approximately 4%. However, road grades vary significantly per road segment. The steepest segments have a gradient around 15%, while other roads are near-flat.

Figure 3.1 is converted into a schematic drawing, which is displayed in Figure 3.2. This drawing is a simplified representation of the mine. It provides an overview of the infrastructure of the mine and the locations of the most important vehicles. Later in this chapter the schematic will be expanded with the operational processes and the rules that guide/constrain these processes. This last schematic will serve as a key input for the model constructed in Chapter 4.

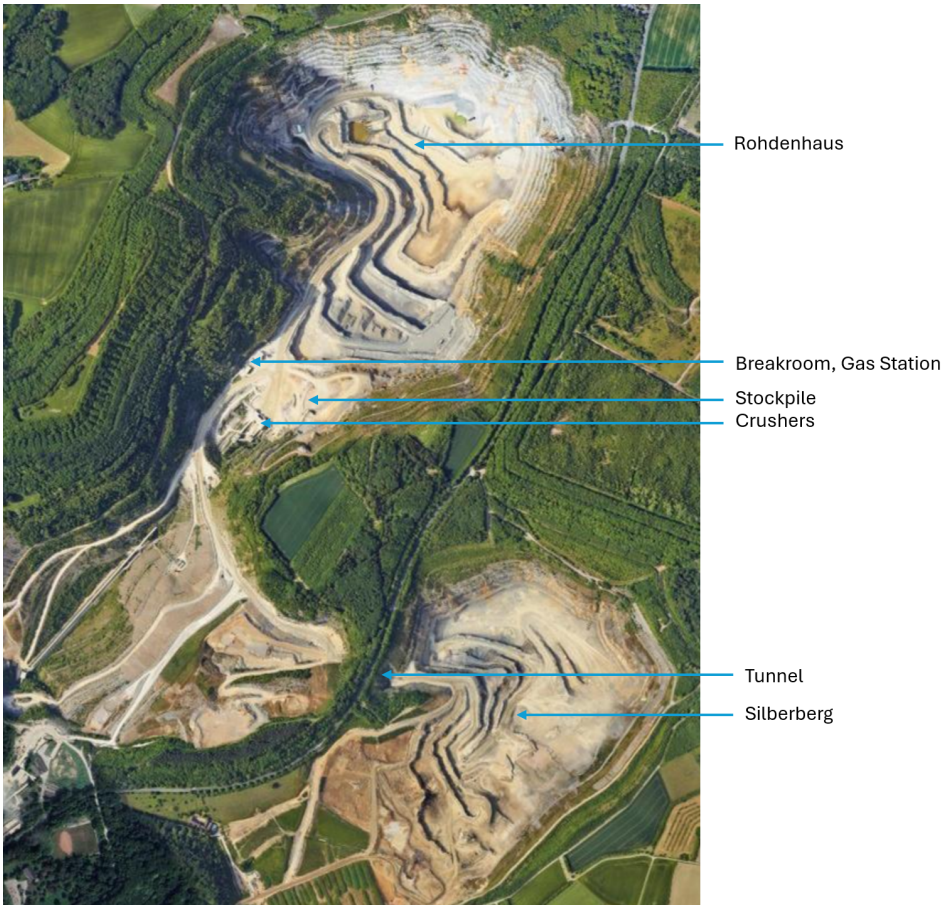

Figure 3.1: Aerial overview of the Flandersbach mine, indicating the two pits, crushers, stockpile, key facilities and the tunnel.

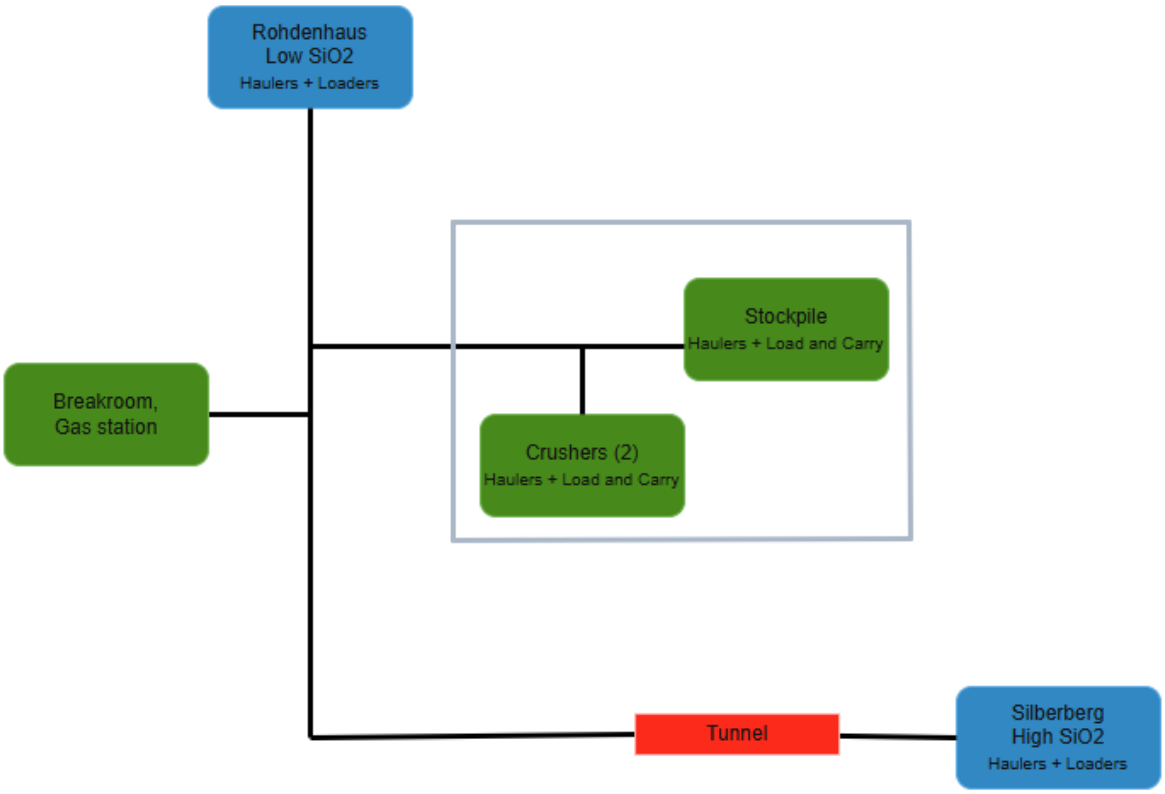

Figure 3.2: Schematic representation of Flandersbach, distinguishing material sources (pits), destinations (dump spots and facilities) and the tunnel.

Figure 3.2 displays destinations in green (crushers, stockpile and facilities), material sources (the two pits) are displayed in blue. The gray box depicts an important part of the mine, hereafter referred to as the main dumping area, as this is the location where most of the material is hauled to. The load and carry vehicle (LC) is confined to this area. Furthermore, within this area most of the interactions between vehicles occur. This is because every hauler drives from one of the pits to the main dumping area and the LC is confined to this area as it drives from the stockpile to the crushers, or performs a cleaning job at the main dumping area. The tunnel is shown as a red rectangle, and is located at the road that connects the Silberberg pit to the remainder of the operation.

3.2. Mine Fleet

The Flandersbach operation employs a wide range of vehicles. However this thesis focuses on the three primary vehicle types: haulers, loaders and LC. Ancillary vehicles are present too including a grader, a water truck, drills, and light vehicles. The light vehicles are 4x4 cars and busses, which are used to get around the mine and transport employees. Occasionally, when the fuel reservoir is empty a large fuel truck enters the mine. These vehicles follow the same safety rules as the primary vehicles, but fall outside the scope of this study and are not modeled.

The haul fleet consists of ten haul trucks, eight Caterpillar haulers and two Komatsus. The sizes and specifications of these trucks are comparable. Four loaders are also utilized. Three of these are Caterpillar wheel loaders, the last one is a Hitachi excavator. The LC is an essential vehicle type in the mine, but it does not correspond to a dedicated machine class as it refers to an operational use of the wheel loaders. These wheel loaders therefore have multiple functionalities. Tables 3.1 and 3.2 list the modeled haulers and loaders respectively, together with their key specifications. These values have guided the modeling process as they indicate the differences and similarities between the vehicles. The asset IDs used per vehicle are reported too, these IDs consist of a prefix which indicates the vehicle type and a number to differentiate between the vehicles of the same type. These names follow the mine's naming convention. In this thesis vehicles are referred to as Hauler 17, Loader 54, etc. However, the original German asset IDs (SKW 17, LG 54) appear in the company documents referenced in this thesis. The loader type in Table 3.2 indicates which loaders can serve as LC, as only the wheel loaders are able to perform the LC task.

Table 3.1: Specifications of the Modeled Haul Truck Fleet

Asset ID	Model	Payload (t)	Max Speed (km/h)	Hoist Up (s)	Hoist Down (s)
SKW 16	Cat 777D	90.1	65.0	15	13
SKW 17	Cat 777D	90.1	65.0	15	13
SKW 28	Cat 777G	90.8	67.1	15	13
SKW 29	Cat 777G	90.8	67.1	15	13
SKW 32	Cat 777G	90.8	67.1	15	13
SKW 36	Cat 777G	90.8	67.1	15	13
SKW 37	Cat 777G	90.8	67.1	15	13
SKW 38	Kom HD 785-7	91.0	65.0	16	14
SKW 39	Kom HD 785-7	91.0	65.0	16	14
SKW 40	Cat 777-07	91.7	65.9	15	13

Data sourced from manufacturer specifications Caterpillar Inc., 2000, 2013, 2022; Komatsu Europe International N.V., 2018.

Table 3.2: Specifications of the Modeled Loader Fleet

Asset ID	Model	Type	Bucket Volume (m^3)
LG 14	Hitachi HCM EX1200	Excavator	6.5
LG 53	Cat 992K	Wheel loader	12.3
LG 54	Cat 992K	Wheel loader	12.3
LG 55	Cat 992	Wheel loader	12.3

Data sourced from manufacturer specifications Caterpillar Inc., n.d. Hitachi Construction Machinery Co., Ltd., 2018; Zeppelin Cat, n.d.

3.3. Production Cycle

Production at Flandersbach, and in mining in general, follows a cyclical process. In this process the material is moved from the benches to the crushers, and sometimes to a stockpile as an intermediary stage. Afterwards haulers travel to the benches again and the cycle is repeated. This cycle can be broken down into several stages.

The cycle begins with the blasting process. First holes are drilled by a drill, then these holes are filled with explosives. After detonation the rock is fragmented such that it can be loaded, transported, and fed to the crushers. At Flandersbach blasting is done in both the Rohdenhaus and the Silberberg pit. Typically, a single blast is conducted each day across the operation.

The loading cycle starts with a loader scooping up the rock with its bucket. Once a hauler arrives the loader lifts its bucket in the air and the hauler positions its tray underneath the bucket, a process herafter reffered to as ranging at a loader. The hauler remains stationary and the loader operator dumps multiple buckets in the hauler's tray. This loading stage ends when the loader operator determines that the hauler payload is sufficient, according to company guidelines outlined in next section and based on the display located at the side of the hauler.

Once loaded, the hauler departs from the load location and travels via the road network to its destination. Possible destinations are a waste dump, a crusher or the stockpile. Because there occurred no dumps at the waste dumps during the simulated weeks these locations were excluded from this research. When the hauler arrives at the main dumping area it dumps its load by lifting its tray. The hauler operator has to choose between one of the crushers or the stockpile, depending on the crusher status and the characteristics of its load. This choice is guided by rules outlined in the next section. After dumping the hauler returns to the same loader and the cycle repeats.

At the processing site, the LC is present as well. This machine has two primary tasks: managing the main dumping area and feeding the crushers. Because haulers dump large volumes of rock at both the stockpile and at the crushers spillage often occurs. These rocks create tire damage risks so regular cleanup by the LC is recquired. Spillage occurs at the crushers by incorrectly positioned hauler trays, or when a load is dumped in a full hopper. At the stockpile spillage occurs when a load is dumped near the crest, because of the height of the stockpile individual rocks may roll beyond the pile boundary.

3.4. Rules and Schedule

The previous section described the material flow throughout the operation. This flow is governed by operational rules and constraints and is structured by a specific schedule. This section first summarizes the rules and constraints, subsequently two distinct schedules for a typical day are presented. All of these elements constitute important inputs for Chapter 4. A summary of these operational rules and constraints is presented in Table 3.3, the two schedules are depicted in Figure 3.4.

3.4.1. Loading

The balance between operational haulers and loaders is critical in mining operations. An imbalance will inevitably lead to idle times for the hauler fleet or the loader fleet. However, achieving a perfect balance is very difficult if not impossible. When there are a lot of loaders compared to haulers the system is undertrucked. In such a configuration the loader fleet experiences unavoidable idle times. At Flandersbach the fleet is set up in an undertrucked configuration. Because of this configuration the loading policy at Flandersbach is to load the haulers as full as possible, while remaining within site payload targets and avoiding systematic overloading to ensure safety and protect the hauler from wearing down too fast. Because hauling capacity is the limiting factor queues are uncommon, and an extra cycle will not cause these queues in most cases. This contrasts with an overtrucked configuration. In such a system the logical policy would be to fill haulers less full. Adding an extra cycle in this setup would result in large queues at the loader, due to the comparatively high amount of haulers present.

To support the loader operator's judgment of the payload every hauler has a screen which displays the number of tons in the tray. This way the loader operator is able to see how full the truck is, as it is impossible to look into the tray from the position of the loader operator. The hauler has to range in such a way that the loader operator is able to see this screen. Once the truck is full enough the loader

operator will signal to the hauler operator that he/she may depart.

3.4.2. Hauling

Currently, the haulers are assigned to a specific loader at the start of a shift. This one loader is the only loader they will travel to. This is a rigid setup, referred to as fixed assignment hereafter. The number of operational haulers and loaders during a shift is based on the production target. This target is set via a negotiation between the plant and the pit supervisor. This negotiation consists of finding a balance between what the plant needs, and what the pit supervisor thinks is possible to haul. This in turn relies on the amount of operators and vehicles available. The number of haulers per loader needs to be determined too. This is done by the pit supervisor, who makes this decision based on his experience. The primary consideration in this decision is the cycle time of the haulers from the load location to the main dumping area. A far away loader will have a longer average cycle time, which means that that loader can service more haulers without queue formations because of the longer travel time.

Haulers must maintain a safety distance of 50 meters from every other vehicle in the mine, and their maximum speed is capped at 30 km/h for safety too. The only time the safety distance rule is relaxed is at the main dumping area, as this rule is not viable here. There are too many vehicles present at this site and haulers have to dump together sometimes as per another rule. The speeds of the haulers is well below the maximum speed in this area, making the removal of this safety distance less of a risk. The roads in the mine are two-way roads, with the exception of the tunnel which is a single-lane road. To prevent collisions overtaking is prohibited everywhere in the mine.

3.4.3. Blasting

Blasts are conducted on most days. When a blast is performed everyone except the blast crew must exit the pit where the blast is performed, due to safety risks.

3.4.4. Dumping

At the main dumping area the important rules the operators need to obey mainly relate to the capacity of the crushers and the grade of the feed. The grade is measured indirectly by examining the SiO_2 content with an XRF which measures the grade of the outgoing crushed rock at the conveyor belt, and displays this percentage at a screen near the crushers. The SiO_2 needs to be kept below 1.7%. The amount of mud needs to be controlled too as this is the main cause of belt failure. To manage these two constraints haulers occasionally need to dump together. Whenever a truck with a very low grade arrives at the crushers this hauler has to wait for a hauler with a high grade. This way the average grade can be managed. The same applies for a muddy hauler. In this case the muddy hauler should wait for a hauler with less mud, to manage the average mud content of the feed. The capacity of the hoppers forms another constraint. The capacity of each hopper is about 300t, but the hopper may only be filled up to 70%. Otherwise, the chance of spillage becomes too big. The current fill percentage is displayed on a screen, this way the operators can instantly see if they may dump or not. If both crushers are too full, or if they are not operational the hauler operators divert to the nearby stockpile.

Table 3.3: Summary of Operational Processes and Rules at Flandersbach

Process	Vehicles Involved	Applicable Rules & Constraints
Blasting	N.A.	All personnel must leave the pit. Often at 11:00 coincides with coffee break to minimize downtime.
Loading	Loaders, Haulers	System is “undertrucked”. Load haulers as full as possible, even if load time increases.
Hauling	Haulers	Speed limit of 30 km/h for all haulers. Safety distance of 50 m separation, and no overtaking, on all haul routes. Rule is lifted at the crusher/stockpile site. Haulers assigned to one specific loader per shift (rigid setup). Check if another vehicle occupies the tunnel before traversing the tunnel.
Dumping (Crusher)	Haulers, LC	Hopper Capacity max 70% fill of the $\sim$300 t hopper. Silica SiO_2 content must be $< 1.7\%$. Muddy trucks must coordinate to dump with a clean truck. Trucks with high SiO_2 content have to dump with a truck with a low content A traffic light indicates the prioritized crusher, but if both crushers$> 70\%$, divert to stockpile.
Stockpile Mgmt.	Load and Carry	Primary focus is feeding the crushers, but the stockpile should be kept tidy
Breaks	All	Haulers must be empty before traveling to the breakroom. Loader operators transported by light vehicles. Load and Carry have a different schedule, breaks 30 min later, to ensure continuous crusher feed.
Refueling	Haulers	Occurs in a specific time window (08:00–09:30). Only one hauler at a time.

The schematic drawing of Figure 3.2 is expanded to include the vehicles and applicable rules present in Flandersbach in 3.3. The rules are placed at the location/vehicle to which they apply. This schematic drawing will be the main guideline for the model presented in Chapter 4.

3.5. Daily Schedule

To interpret operational behavior at Flandersbach, it is necessary to understand the timing of key activities and interruptions over a typical day. The schedule may vary due to the dynamics within the mine, however under normal circumstances it follows the structure shown in Figure 3.4.

The mine operates five or six days a week. A regular working day has 15.5 operational hours, whereas Fridays comprise 13 operational hours. Occasionally the mine is operational on Saturdays too. These days are always split in two shifts, hereafter referred to as the early shift and the late shift.

The early shift starts at 6:00. At this time the hauler operators mobilize their haulers and start driving to the loaders. The crushers are turned on after 15 minutes, reflecting the time it takes for the first loaded haulers to reach the main dumping area. Work continues until 11:00, when hauler and loader operators take their coffee breaks. This time is chosen on purpose as this coincides with standard blast time. This way the production loss caused by the blasting safety rule to leave the pit during a blast is

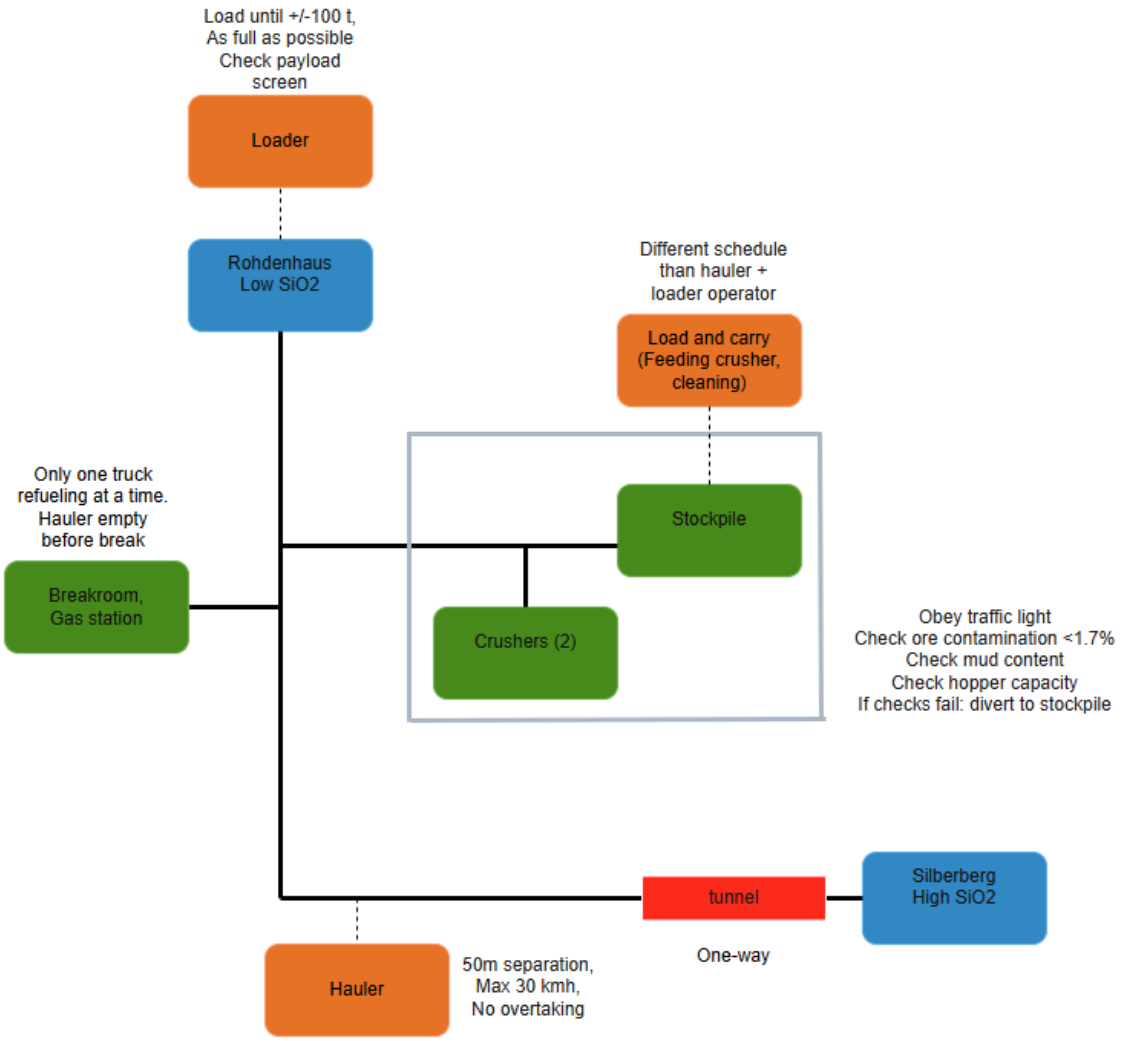

Figure 3.3: Schematic map of Flandersbach with vehicles and rules.

minimized. After half an hour the operators turn back to their respective tasks. Hauler operators may only go to the break room after they have emptied their truck. Loader operators do not use their loader for travel, but are picked up by a light vehicle.

The early shift ends at 13:30. Again haulers have to be emptied before traveling to the break room and the loader operators are collected by a light vehicle. The employees are then brought from the break room to the entrance and the operators of the next shift take their place. This takes 30 minutes in total. The LC operator follows a slightly offset schedule and stops half an hour later. Therefore during the shift change the crushers are still being fed, and they do not lose their purpose during the shift change. The same holds for the breaks.

The mine closes at 21:30 on regular days, and on 19:30 on Fridays. The timing of the stoppages on Friday is different but the duration is not. The differences between regular days and Fridays can be seen from Figure 3.4.

Refueling is done in a specific time window. The hauler operators refuel their vehicle between 8:00 and 9:30. There is one fast refuel hose, which fills a vehicle within about seven minutes. Therefore, only one hauler may refuel at a time. This keeps the continuity of the operations as big as possible, and minimizes the time spent tanking because only the fast refuel hose is used.

3.6. Variability and Flexibility in Flandersbach

As introduced in Chapter 2, mines are influenced by variability of processes and the system's flexibility to absorb or counteract this variability. In practice both concepts arise at many points in the mine. the examples below serve as concrete real life illustrations for both effects.

Operational variability can propagate through the tightly coupled mining environment. One source is heterogeneity in operator behavior. Under similar conditions, a cautious or inexperienced operator may drive at a slower speed than a more experienced operator. In Flandersbach, where overtaking is not allowed, such speed differences can lead to bunching, where a set of haulers is stuck behind the first hauler with the lowest speed. A second source is variability in service processes at shared resources, like crushers and loaders. Haulers may arrive at a loader with headways on the order of one loading time (or more), which can support stable operation when service times remain within a typical range. However, a single unusually long loading event can reduce effective headway and trigger queue formation. Once a queue forms, the system may remain in a congested regime for an extended period because subsequent haulers inherit delays. These examples illustrate how minor deviations at the vehicle or interaction level can scale into system-level effects such as increased waiting times, irregular crusher feed, and short-term production losses.

A central flexibility mechanism at Flandersbach is the availability of two crushers, which provides redundancy in processing capability. To illustrate the effects of flexibility, a simplified case is presented below where crusher failures are independent and each crusher is down with probability:

$$p(\mathsf{fail}) = 0.1.$$

The probability that both crushers are down:

$$p(\mathsf{no\_crushers\_available}) = p(\mathsf{fail})^2 = 0.01,$$

implying a substantially higher probability that at least one crusher remains available, compared to a situation with one crusher and thus no redundancy, where this chance would be 10%.

$$1 - p(\mathsf{no\_crushers\_available}) = 1 - p(\mathsf{fail})^2 = 0.99$$

While this calculation is intentionally simplified, it demonstrates the general mechanism: redundancy reduces the risk of complete loss of crushing capability and therefore increases the system's ability to continue operating under disturbances. A second form of flexibility is the availability of reserve capacity in mobile equipment, which can partially compensate when a unit breaks down. More generally, the available flexibility determines whether variability is absorbed locally with limited impact, or amplified into sustained queues, under-utilization of critical assets, and increased downtime.

Together, these Flandersbach examples motivate the modeling and experimental emphasis of this thesis. The subsequent chapters quantify how dispatching choices and autonomy-related interventions change the system's sensitivity to variability and the extent to which available flexibility can be leveraged to maintain throughput and compliance under realistic disturbances.

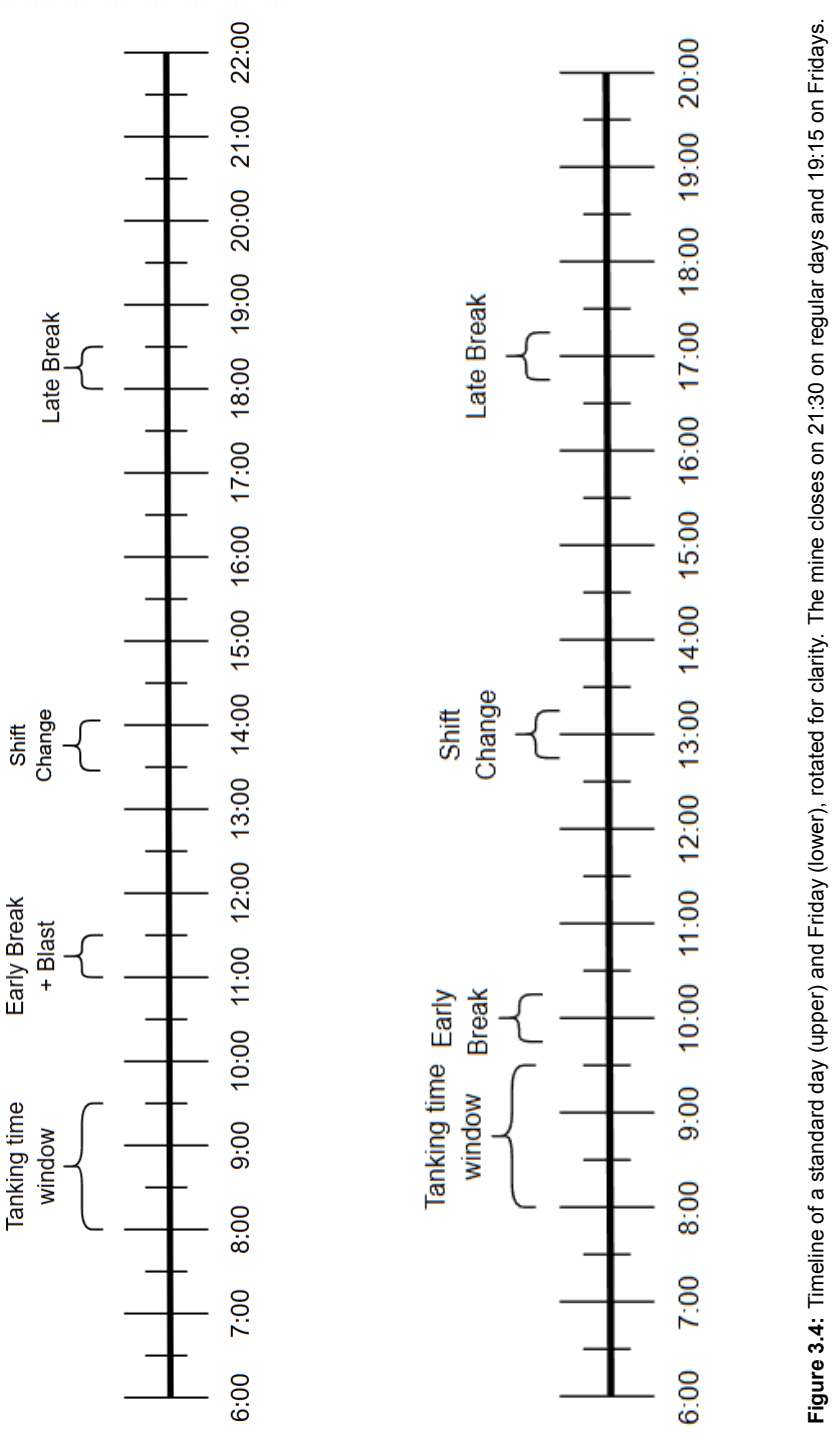

Figure 3.4: Timeline of a standard day (upper) and Friday (lower), rotated for clarity. The mine closes on 21:30 on regular days and 19:15 on Fridays.

4

Methodology

This chapter outlines the methodological approach taken to address the research questions. The structure of this chapter follows the order of the steps taken in the development of the model. First, the data analysis will be described, which renders the inputs for the Discrete Event Simulation via empirical distributions derived from trimmed telemetry data. Then, the architectural setup of the DES will be explained. After this, the operationalization of the rules and constraints in the mine will be highlighted. The methods and metrics for the experiments will be described in the subsequent sections. Finally, the visualization of the model is presented which serves to verify the model workings are implemented correctly and allows for a better understanding of the model behavior.

4.1. Data Analysis and Operationalization

The DES is parameterized using empirical distributions from processed telemetry data. The telemetry data has been recorded via onboard measuring systems, and retrieved via the Caterpillar VisionLink platform, which provides the data for the hauling, loading and LC activities. The VisionLink records are transformed into DES inputs through a preprocessing pipeline that screens data quality, fills in missing data through matching processes and constructs empirical distributions for the stochastic processes within the DES.

The preprocessing pipeline filters values that are not physically plausible and appends missing operational identifiers (bench and assignment attributes) through matching to internal company documents. The resulting pipeline yields cleaned datasets that serve as direct inputs and validation targets, for the DES. Table 4.1 summarizes the data sources and extracted variables used per vehicle type and indicates how these data sources are used within the model. Record counts are reported in Chapter 5 for the telemetry-derived vehicle datasets.

The observed differences between vehicles were compared against the manufacturer specifications from Table 3.2 and 3.1. This comparison serves as a plausibility check to guide the selection of the most reliable datasets, and consequently construct the most reliable empirical input distributions. In addition to record matching, internal company documents are used to derive shift-level production totals for the simulated weeks. On-site observations and discussions with mine personnel were used to formalize operational rules not observable in telemetry data including no-overtaking behavior, dumping and coordination practices at the crushers, main dumping zone headway exceptions, and break/shift-change behavior. The company schedule and production reports excerpts used in this study are provided in Appendix F, accompanied with an explanation regarding the assumptions and the methods used to interpret these documents.

4.1.1. Preprocessing of Telemetry Data

The plots used to assess raw telemetry data quality are provided in the appendices (haulers: A, loaders : B, and LC: C), accompanied by vehicle specific notes describing their data availability and anomalies in their data in Chapter 5.

Table 4.1: Overview of data sources and extracted variables used for simulation parameterization.

Vehicle type / subsystem	Primary data source	Extracted variables	Use within the model	Preprocessing and inclusion criteria
Haul trucks	On-board telemetry (cycle-level)	Travel time and distance (derived speed), dumping time, payload	Dump time is used as a stochastic input. Speed and payload are used as calibration and validation references.	Removal of physically implausible records. Completion of missing identifiers via matching to internal documents. Derivation of speed from distance and travel time.
Loading equipment	On-board telemetry (bucket- and/or cycle-level)	Bucket payload, bucket cycle time, loading cycle characteristics	Bucket payload and cycle time are used as stochastic inputs for truck loading.	Exclusion of inconsistent bucket sequences and implausible loads. Selection of representative bucket cycles. Removal of erroneous extremes.
Load-and-carry	On-board telemetry (cycle-level)	Payload, digging time, dumping time, loaded travel time, empty travel time	The listed variables are used as stochastic inputs for the load-and-carry process.	Selection of cycles corresponding to crusher-feeding routes. Conversion of logged durations to consistent units. Removal of invalid durations and implausible payloads.
Crushers and dumping facilities	Site documentation and operational confirmation	Buffer capacity, processing rate, number of dumping positions, operating constraints	Service capacities and constraints are implemented deterministically. Availability behaviour is defined by a downtime generator.	Parameters fixed to observed operating conditions. Coordination constraints confirmed through site discussions.
Shift schedule and production context	Internal planning schedule and production reporting for simulated weeks	Bench–loader–truck assignments, planned tonnage allocations, crusher vs. stockpile totals	Inputs to scheduling logic, production targets, and reporting totals.	Reconciliation of identifiers with telemetry. Correction of reporting artefacts in operating-hour aggregation.
Road network and topology	Company file with string data	Node coordinates, segment lengths, segment grades, terminal locations	Routing and travel-time computation. Grade-dependent speed modeling.	Connectivity checks and manual completion of missing terminal locations where required.

Haulers

Hauler travel speed is not provided directly in the telemetry data, but constitutes an important input in the DES. Therefore, it is derived from recorded travel time and travel distance in loaded and empty states. Travel distance is missing for a substantial amount of cycles, motivating an imputation step

prior to this speed calculation.

The missing travel distances are imputed using the mean distance of similar telemetry records. Similarity, in this case, is defined by longitude and latitude proximity, identical destination, and the same calendar day. The longitude and latitude were considered to be similar if their difference was +/- 0.00045 degrees, which amounts to about 50 meters difference in latitude and about 35 meters difference in longitude. The calendar day is an important matching criterion as the location of benches, and the road network itself changes over time altering the travel distance associated with the longitude and latitude coordinates substantially. The destinations needs to be the same, because the travel distance will be altered when the haulers travel to other destinations.

In addition to distance imputation, a second matching procedure assigns a bench to each telemetry record. Bench-level differentiation is required for reliable travel-speed calibration, because travel speeds vary due to differences in road geometry and depth between benches.

Bench matching is performed by linking telemetry cycles to the company schedule, using matching criteria which consist of: calendar day, shift, loader ID, hauler ID, pit name, and destination. This combination was chosen because the mine operates with a fixed assignment policy, and loaders typically remain at a single bench during their shifts. Consequently, when a loader in the telemetry data matches a loader in the company schedule on the same day and shift, the bench recorded in the schedule can be assigned to the corresponding haul cycles. Pit was included to avoid incorrect matches between benches with the same bench number in the two pits, and destination was used to keep only productive cycles relevant for travel-speed analysis. The resulting assignments were verified through a spatial plausibility check by plotting the telemetry coordinates on an aerial image of the mine and assessing wether the bench labels formed coherent spatial clusters. This figure is included in Appendix A in Figure A.1. Cases in which loaders or haulers appeared at multiple benches within a shift are treated as non-unique and are removed.

Lastly, telemetry variables used by the DES are converted to consistent units and compiled into empirical sample sets. These sample sets represent the stochastic inputs used to model dump time and payload transfer in the simulation. Extreme values are excluded through deterministic screening in preprocessing, while the sampling-specific trimming strategy applied during simulation input construction is described in the vehicle-process section.

Loaders

The loading process is parameterized with bucket payload and bucket cycle time as inputs. Therefore, preprocessing focuses on the identification of coherent loading cycles to construct empirical distributions for these two parameters. A coherent loading cycle is defined as a loading cycle with bucket-sequence indexes decreasing stepwise by one for each subsequent bucket until the end of the load cycle.

Records indicating implausible bucket payloads or cycle times are excluded. Only coherent loading cycles are kept in the data analysis, cycles where the bucket-sequence pattern is corrupted are excluded from the data analysis too.

The empirical inputs are constructed to represent steady-state loading. Therefore, bucket cycle times belonging to the first load of a cycle (bucket sequence of 1) are excluded from the input. These cycles include ranging and coordination times that would bias the cycle time distributions. Bucket payload distributions are restricted to bucket sequence indices belonging to full buckets (bucket sequence <6), as the DES does not model partially filled buckets. Retaining the bucket payloads with higher sequence numbers would systematically reduce the sampled bucket payloads.

Load and Carry

The LC process is parameterized using empirical distributions for payload, digging time, dumping time, and travel times in loaded and empty states. Telemetry fields required for this parameterization are incomplete for the unit operating during the simulated weeks, empirical inputs are constructed from telemetry collected for comparable equipment performing the same operational role.

LC also performs auxiliary tasks in addition to crusher feeding, such as cleaning activities in the main dumping area, which reduces the feed rate of the LC. To quantify this, 'Other Activities' are defined as

any gap exceeding one minute between recorded telemetry cycles. This metric tracks the proportion and variability of time spent outside primary operations, providing essential context for interpreting LC results in Chapter 5.

4.2. Discrete Event Simulation

This section describes the construction of the DES environment used in this study. The model is parameterized using the processed telemetry data described in the previous section. Quantitative validation against measured site performance is reported in Chapter 5. This includes component-level comparisons, covering both direct input variables and downstream outcomes used for validation, and system-level comparisons, which focus on production over the simulated weeks. This section describes the simulation environment used for the validation and evaluation of all experiments, and it is structured in two parts. First, the architecture and core components of the DES are described. Second, the implementation of the operational logic and constraints is outlined. To support the reading of the methodology chapter, a schematic overview of the DES architecture is provided in Appendix E in Figure E.1. The figure indicates where the main submodels act in the Flandersbach system and clarifies the hierarchy between operational submodels and supervisory layers such as dispatch and schedule logic.

4.2.1. Architecture and Core Components

The DES is constructed as a digital representation of the Flandersbach mine to simulate the baseline operation and experimental scenarios introduced later in this chapter, that serve to answer the research questions of this thesis. The model makes use of the SimPy library, which is specifically designed for DES through process-based event scheduling, time-ordered resource contention, and explicit modeling of queues and buffers.

The haul road network is modeled as a directed graph, with the exception of the tunnel which is a single lane road, using NetworkX. The graph is generated by converting string data of the mine from company files, to create different nodes with an X,Y,Z coordinate. In between these nodes edges were made, and a physical length and grade percentages were assigned to these edges. The terminal nodes of this network were converted to load locations with labels corresponding to their real world counterparts. The string data did not include the coordinates for the crushers, stockpile, petrol station and break room. Therefore, these facilities were manually added to the network at their corresponding locations, which can be seen in Figure 3.1. A 2D visualization of the resulting network is presented in Chapter 5 to support model verification.

The modeled infrastructure includes two crushers, a stockpile, a break room, and a fuel station. Each crusher is parameterized by a processing rate of 1000 t/h and a hopper buffer capacity of 300 t. These parameter values reflect the operating configuration during both simulated weeks. Dumping at each crusher is modeled by two dumping bays, while each loader provides a single loading position. The stockpile is modeled as an unbounded buffer and is initialized with 2000 t to enable immediate LC operation at the start of the simulation. The break room and fuel station are represented at a shared location, consistent with the site layout as can be seen from Figure 3.1.

4.2.2. Operational Logic and Constraints

This subsection describes how the operational rules and constraints identified for the Flandersbach mine are implemented in the DES. Table 4.2 summarizes the translation from the operational requirements presented in Chapter 3 to the corresponding simulation logic. The remainder of this subsection details the transportation model, vehicle-process representations, safety and capacity rules, and the schedule logic used during the simulated weeks.

Transportation

Hauler transportation in the DES is modeled as traversal in a mostly bidirectional road network. This section describes the derivation of the achieved speed as a function of road grade, load state and operator profile with the associated stochastic variations. After this, constraints regarding safety distance and tunnel capacity are described.

Every edge, is assigned a grade and length based on the starting and ending node coordinates of the edge. The velocity computation then starts with the setting of a base speed, which depends on the pit

Table 4.2: Translation of operational constraints, rules, and values to simulation logic.

Operational constraint (Chapter 3)	Simulation implementation (Chapter 4)
Haul road network	Represented as a directed graph in `NetworkX`. Nodes store 3D coordinates and edges represent road segments with length and grade attributes.
Speed limits	A global speed limit of 30 km/h is enforced on haul-road segments. A speed cap of 5 km/h is applied near loading and dumping locations.
Driver variability	Operator variability is represented using an aggressiveness factor governed by a bounded random walk, which alters the baseline segment speed.
Safety distance (50 m)	Road segments impose a capacity constraint that limits the number of concurrent vehicles on an edge. Capacity is derived from edge length using a 50 m spacing proxy.
Tunnel bottleneck	The tunnel is modeled as a shared resource with capacity one. Vehicles must acquire the resource before entering to prevent opposing direction conflicts.
Loading policy (load full)	Loading is modeled as repeated bucket cycles until the hauler payload exceeds a threshold of 87 t.
Loader variability	Loading durations and bucket payloads are sampled from empirical distributions. Excavator loading is represented using an additional full-cycle time penalty to reflect lower bucket capacity.
Ranging (maneuvering)	Maneuvering time is represented through fixed time allowances. Different allowances are applied for loading and dumping.
Dumping logic	Vehicles check crusher buffer status prior to dumping. When the buffer exceeds the diversion threshold (62%), haul trucks divert to the stockpile.
LC transportation	LC travel is represented using aggregated travel-time sampling rather than explicit graph traversal, reflecting the short route between stockpile and crushers.
Crusher reliability	Crusher availability is primarily modeled using a stochastic downtime generator with Markov-type switching. The mean downtime duration is 30 minutes.
Refueling	Each haul truck refuels once per day during a fixed time window (08:00–09:30). Refueling duration is modeled as 7 minutes.
Shift schedule	Shift logic is implemented using time-indexed schedules. The model includes Friday-specific operating hours (operations cease at 19:15) and early-break decision logic based on a 14-minute look-ahead.

the hauler is traveling from/to due to significant mean speed differences between the two pits. In the speed calculation the speed is not altered when the hauler is cornering, and the calculation neglects acceleration/deceleration. The acceleration/deceleration has been approximated by a limited speed at terminal edges as will be explained later. The pit-specific baselines reflect differences in the road network of the two pits, related to the total length and amount of corners of the roads from the two pits. The two pits share no roads, with the exception of the road leading to the crushers/stockpile/break room but these edges have assigned fixed speeds. Bench specific speeds were deemed impractical, as the benches in the pits share a lot of roads. This would have resulted in different speeds for haulers traveling to a different bench in the same pit on the same edge, which is unrealistic in most cases.

A second split is the differentiation between loaded and empty states for the mean speeds. The last parameter influencing the base speed is the grade. The formula used within the model is described in equation 4.1, and the per pit/state parameters are shown in Table 4.3. The calibration factor has been introduced to account for differences in modeled speeds per bench and the speeds per bench found from the empirical data.

$$\mathbf{v}_{\text{pit, state}} = \mathbf{v}_{\text{flat,pit,state}} \cdot (A + B \cdot \text{grade}_{\%}) \tag{4.1}$$

Where:

- $v_{\text{pit,state}}$: Modeled segment speed (km/h) for a given pit and load state.
- $v_{\text{flat,pit,state}}$: Baseline speed on nominally flat grade (km/h) for the same pit and load state.
- $\text{grade}_{\%}$: Signed road grade expressed in percent.
- A: Dimensionless intercept parameter.
- B: Grade sensitivity coefficient (per percent).

The constants A and B, along with the pit-specific and calibrated baseline velocities, are detailed below and summarized in Table 4.3. The grade sensitivity coefficient is set to yield realistic uphill and downhill speed adjustments. Its influence is subsequently evaluated in the sensitivity analysis. The relation between grade percentage and speed reduction/increase is modeled linearly.

Table 4.3: Final calibrated speed parameters for the DES transportation model.

Pit / State	$\mathbf{V}_{\text{Flat}}$ Base (km/h)	Calibration factor	Decay coeff. ($\mathbf{B}$)	Final calibrated $\mathbf{V}_{\text{Flat}}$ (km/h)
Loaded				
Rodhenhaus	10.49	$\times 1.40$	-0.030	$\mathbf{14.69}$
Silberberg	13.95	$\times 1.20$	-0.030	$\mathbf{16.74}$
Empty				
Rodhenhaus	14.12	$\times 1.20$	-0.020	$\mathbf{16.94}$
Silberberg	16.53	$\times 1.30$	-0.020	$\mathbf{21.49}$

Using Equation 4.1, a segment speed is computed for each edge and when needed capped at 30 km/h to comply with the mine's speed limits. To avoid unrealistic speeds near loading and dumping locations, a terminal-area speed cap of 5 km/h is applied on designated terminal edges. This approximates acceleration and deceleration effects that are not modeled in the DES.

The speeds derived from Equation 4.1 are deterministic. To represent operator-driven variability in driving behavior, the model adopts the aggressiveness-factor approach proposed by Parreira (Parreira, 2013). It functions via two attributes assigned to every hauler operator: an aggressiveness factor and a stability factor. The aggressiveness factor shows if the driver prefers to drive fast, in case of a high aggressiveness factor, slow or normal. The range this factor takes is set via the stability factor. The values used have been directly adopted from Parreira's work and are shown in Table 4.4 Parreira, 2013.

Table 4.4: Aggressiveness Factor

Aggressiveness	Stability Factor (Range: $\mathbf{AF}_{\min}$ to $\mathbf{AF}_{\max}$)		
	Very Stable	**Limited Change**	**Variable**
Passive	-1.00 to -0.80	-1.00 to -0.50	-1.00 to -0.20
Normal	-0.10 to $+0.10$	-0.25 to $+0.25$	-0.40 to $+0.40$
Aggressive	$+0.80$ to $+1.00$	$+0.50$ to $+1.00$	$+0.20$ to $+1.00$

The aggressiveness factor itself is not a static value. Rather it is computed every time step, governed by a bounded random walk. This ensures the aggressiveness factor never exceeds the values from

Table 4.4. The value is calculated via Equation 4.2, it depends on the aggressiveness factor from the previous time step, and a random variation.

$$\mathbf{AF}(t) = \mathbf{AF}(t-1) + \Delta af \quad \text{subject to } \mathbf{AF_{min}} \leq \mathbf{AF}(t) \leq \mathbf{AF_{max}} \tag{4.2}$$

Where:

- $\mathbf{AF}(t)$: Aggressiveness factor at simulation time step t (dimensionless).
- $\mathbf{AF}(t-1)$: Aggressiveness factor at the previous time step (dimensionless).
- Δaf: Stochastic step term applied at each time step (dimensionless).
- $\mathbf{AF_{min}}$: Lower bound for the aggressiveness factor (dimensionless).
- $\mathbf{AF_{max}}$: Upper bound for the aggressiveness factor (dimensionless).

Once this aggressiveness factor is determined it is used to model the hauler's actual driving speed. This is done by multiplying the speed from the edge for the state that hauler is in via equation 4.3. The Kappa value has been set to 0.15, by comparing achieved speed distributions in the simulation to the speed calculated from the raw data. This comparison is reported in Chapter 5.

$$\mathbf{v}_{\text{drive}} = \mathbf{v}_{\text{pit, state}} \cdot (1 + \kappa \cdot \mathbf{AF}) \tag{4.3}$$

Where:

- $\mathbf{v}_{\text{drive}}$: Realized driving speed after applying behavioural variability (km/h).
- $\mathbf{v}_{\text{pit,state}}$: Deterministic segment speed for the relevant pit and load state (km/h).
- $\mathbf{AF}$: Aggressiveness factor (dimensionless).
- κ: Scaling parameter controlling the magnitude of behavioural variability (dimensionless).

Vehicle Processes and Modeling

As detailed in Chapter 3, each vehicle follows an operational cycle, which is intertwined for haulers and loaders. The DES uses empirical sample sets constructed from telemetry-derived variables described in Section 4.1. Prior to sampling, all empirical inputs are trimmed to remove outliers by excluding values outside the 2.5th and 97.5th percentiles.

Hauler Process

The only empirical dump time used for haulers is the dump time, which is modeled using an empirical distribution derived from telemetry records. Because truck specifications are similar and telemetry coverage is uneven across vehicles, the dump time input is taken from the subset of records that exhibits plausible behavior compared to manufacturer specifications. A fixed ranging time is included to represent positioning time at the crusher and loader.

Loader Process

The loader process is modeled as repeated bucket cycles that continue until a payload threshold is reached, reflecting the Flandersbach policy to load haulers as full as possible. The DES does not model partial scoops, and the threshold is calibrated to reproduce the payload distribution for haulers with reliable payload telemetry data. The loading activity takes the processed empirical distributions for payload and cycle time as its inputs.

Load and Carry Process

The LC is modeled to exclusively feed the crushers, in contrast to reality where it performs additional auxiliary tasks. The road between the crusher and the stockpile is extremely short, and the safety distance is lifted at this particular spot within the DES. Instead of a speed calculation on this single edge the LC movement has been modeled as a time-based stochastic. This means the LC does not traverse any edges, rather it teleports from the stockpile to the crushers and back. The time the teleportation takes is taken from the distributions of its loaded travel time, and empty travel time, for cycles going to/from the crusher. The other inputs of the LC process include the dig time, dump time and payload per bucket. These are sampled from the preprocessed empirical distributions.

Safety, Scheduling and Capacity Rules

This section describes the rules implemented in the model that reflect the operational policies of the site, including safety distances, facility capacities and schedule logic.

Safety Distance

The speed calculated via equation 4.1 and 4.3 is not the speed haulers necessarily achieve when traversing the network, due to the safety rules of the mine. Each edge has an assigned capacity, which is defined as the maximum amount of haulers that may traverse the edge at the same time. This capacity acts as a proxy for the safety distance of 50 meters, and is computed via formula 4.4.

$$\text{Capacity}_{\text{Edge}} = \max\left(1, \text{round}\left(\frac{\text{Length}_{\text{m}}}{50\ \text{m}}\right)\right) \tag{4.4}$$

Where:

- $\text{Capacity}_{\text{Edge}}$: Maximum number of haul trucks allowed concurrently on an edge (dimensionless).
- Length_{m}: Edge length in meters (m).

An edge of 50 meters for example has a capacity of one. This means that in this scenario two haulers can not travel in the same direction on this edge. On short edges the rounded value could result in a capacity of zero, which is undesired as no haulers can traverse that edge. Thus, the minimal capacity has been set to one. This proxy has been used because calculating distance between vehicles at every time step is difficult and computationally expensive. The result of the rounding and the minimum value of one is that, at times, the safety distance requirement may not be met and that two haulers are separated by less than 50 meters from each other. However, as the network is long, these haulers meet an edge where they do have to wait, as the calculation is better suited for longer edges. The rounding may result in a lower than actual capacity too. When the edge length is 74 meters two haulers could traverse that edge in reality, however in the simulation this is not possible as per formula 4.4.

This safety distance rule has been lifted at the terminal edges, which are defined as: edges leading to a load location, a dump location, the break room/ petrol station. This way queue formation is possible at these locations, which is a necessary requirement when evaluating the system. In reality this safety distance rule is lifted at these spots too, due to low speeds and overall infeasibility of the rule. The maximum speed of these terminal edges has been set to 5km/h.

A special case is the edge which represents the tunnel, as this edge does not allow two haulers traveling in opposite direction. In the DES this has been modeled by setting the capacity of the tunnel to one. Before entering this segment a hauler has to request this shared capacity, preventing two haulers from entering the tunnel simultaneously.

Crushers and Dumping

During the two simulated weeks, crusher availability was approximately 90% based on an interview with the mine manager. Detailed timestamps and durations of individual downtime events were not available. Therefore, three distinct downtime modes were implemented within the DES that reproduce the overall availability. The three modes include: 90/10, 93 minute block , Markov 30 minute average. The 90/10 mode has a crusher up-time of 90 minutes, afterwards the crusher is down for 10 minutes. Consequently both crushers are down at the same time in this mode. The 93 minute block puts the crushers down for 93 minutes, at a random moment during each day. Both these failure modes ensure a total up-time of 90% of the total time. The last mode, a Markov-type generator with a mean downtime duration of 30 minutes, produces stochastic downtime patterns and targets approximately 90% availability. Because outages are sampled randomly, realized availability over a given day or week can be slightly higher or lower than this target. The mean of the down time duration is 30 minutes, and the DES checks every ten minutes if the crusher will go up or down. In contrast to the other two modes the crushers have independent downtime calculations. Therefore, enabling the possibility of just one crusher being up. The Markov mode is used for the main analyses and comparisons in this thesis, because, due to the stochasticity of the mode and crusher down times in reality, this mode is deemed the most realistic representation of crusher failure. The 90/10 mode is deemed the most unrealistic mode, this type of behavior is not probable as crusher failure is stochastic. The 93 minute block, or

any other duration, is usable for different analyses. For example during scheduled maintenance of the crusher(s)

The hoppers have a capacity of 300 tons. In reality there may be a small pile of rock on the hoppers, but in the SimPy library this is impossible. Therefore the 70% rule has been altered to a conservative 62% rule, to prevent dumps that exceed the limit of the hopper capacity (The maximum load in simulation does not cause spillage with this percentage). Both the LC and the haulers check the current hopper status before they dump. If the hoppers are full the haulers divert to the stockpile, the load and carry remains in the queue until the hoppers are below 62%. Material quality constraints such as grade and contamination are outside the scope of this study and are therefore not modeled. Consequently, haulers do not wait at the crushers for high grade/clean haulers.

Timeline

In this section the modeling of the timeline summarized in Figure 3.4 is explained. The mine operates from 6:00 until 21:30 on regular days, on Fridays the operation comes to a halt at 19:15. The schedule has been split, as in real life, on a shift level to incorporate the different quantities and load locations. This structure supports reporting and comparison to company production records on a shift level. The breaks, and shift changes have been implemented at their respective times too and follow a certain logic.

The break logic allows trucks to proceed to the break room early, and ensures a hauler is empty before leaving the production cycle. After each dump event, the truck checks whether the next scheduled break or shift change occurs within a certain time window. When this time is not incorporated haulers traveling long cycles have the chance to not reach the break room or spend too little time in the break room. In reality hauler operators check the time it roughly takes to complete the next cycle, if this time compromises their break time too much they go to the break room early. The early break time window also yields an overnight break duration consistent with the time the mine is closed. The optimal duration for the early break time window was found via multiple iterations and is presented in Chapter 5.

Refueling is done once per day per hauler. Every hauler has to refuel and checks if the petrol station is empty between 8:00 and 9:30. The check is done after dumping, therefore a loaded hauler will never go to the petrol station, which is in line with the company policy. The time refueling takes was not recorded but has been set to a value of 7 minutes. This duration was confirmed during discussions with the mine personnel.

4.3. Current Setup, Sensitivity Analysis and Experimental Design

This section translates the objectives from Chapter 1 into executable policies and scenarios for the DES. It first defines the baseline configuration which serves as the primary benchmark, and is used to validate the model. It then describes the experimental scenarios, including dynamic dispatch, autonomous variants of the haul truck, loader, and LC, and three additional scenario variants. The sensitivity analysis is reported immediately after the baseline description, before the experiments are introduced.

To ensure statistical rigor and comparability across simulations, Common Random Numbers (CRN) are used throughout the study. CRN reduces variance in comparative analyses by exposing alternative policies or parameter settings to the same stochastic background. In practice, the run number is used as the seed in both the tuning phase and the experiment phase. During tuning, this means that each candidate set of heuristic weights is evaluated against the same sequence of seeds, while during the experiment phase the same run number is used across the Baseline and all alternative scenarios. As a result, both phases are evaluated under identical stochastic realizations, including crusher failures, empirical draws from the input distributions, and driver behavior profiles. This ensures that observed differences in objective value or scenario outcome can be attributed to the modeled policy or intervention rather than to random variation.

4.3.1. Current Setup

The current setup models the haulers by making use of their cropped dump time data. Each shift assigns haulers to loaders via a fixed assignment policy, preventing haulers from traveling to different loaders during their shift. The assignment of haulers to loaders has been done in accordance to the

actual assignment, obtained from the full historical schedule. A one-day schedule excerpt used to construct the assignments and locations is provided in Appendix F for illustration, accompanied with a brief explanation of how the schedule is interpreted and how assignment decisions are derived from it. Within-shift loader relocations rarely occur in the modeled weeks. When they do occur, the loader is assigned to the bench that accounts for the largest share of its recorded tonnage during that specific shift. In addition, operator driving behavior is initialized by assigning each driver an aggressiveness profile drawn from Table 4.4, with equal sampling probability across all profile categories.

4.3.2. Sensitivity Analysis

The sensitivity analysis evaluates the robustness of the DES by varying the parameters that are not derived from telemetry and therefore based on interviews with personnel, calibrated assumptions or model design choices. Additionally, the extent of the tail trimming of the empirical distributions is varied in this analysis. The outcomes of and values taken for this analysis are reported in Chapter 5. Table 4.5 provides an overview of the parameters included in the sensitivity analysis, and links each parameter to the corresponding model component, building on the sources listed in Table 4.1 and the rule implementation summary listed in Table 4.2.

The sensitivity scenarios varied one parameter, while all others were kept constant. CRN is applied across the sensitivity analysis to ensure differences in outcomes can be attributed to the parameter change rather than stochastic event variation. Each different parameter setting is simulated over ten independent runs. The trimming bounds constitute the only global modification because they affect all empirical sampling inputs simultaneously.

Sensitivity factors are grouped by model mechanism within Table 4.5. Crusher-system sensitivity covers processing capacity, hopper buffering, the acceptance threshold that governs diversion behavior, and the two alternative downtime modes. Transportation sensitivity evaluates baseline speed levels, grade sensitivity, and the magnitude of operator variability. Loading and dumping sensitivity covers the payload-based stopping criterion used to represent the load as full as possible policy, ranging time at loaders and at crushers, and the additional delay applied to represent the excavator's longer effective loading cycle. Input robustness is evaluated by varying the percentile bounds used to trim empirical sampling inputs. Tested settings and resulting effects are reported in Chapter 5.

Table 4.5: Sensitivity analysis overview. The table lists parameters varied in one-factor-at-a-time sensitivity scenarios around the baseline configuration. Tested settings and resulting effects are reported in Chapter 5.

Parameter	Source type	Related model component
Crusher-system sensitivity		
Crusher processing rate (t/h)	operational confirmation	Table 4.2 (Crushers and dumping facilities)
Crusher buffer capacity (t)	operational confirmation	Table 4.2 (Crushers and dumping facilities)
Crusher acceptance threshold (fraction)	calibrated assumption	Table 4.2 (Dumping logic)
Crusher downtime mode (deterministic, Markov-type, single daily block)	model design choice	Table 4.2 (Crusher reliability)
Transportation sensitivity		
Global speed multiplier (mean-speed scaling)	calibrated assumption	Transportation model (4.2.2)
Grade sensitivity multiplier (scales grade coefficient B)	model design choice	Equation 4.1
Operator variability scaling κ	calibrated assumption	Equation 4.3
Loading and dumping sensitivity		
Hauler payload threshold for load-full stopping (t)	calibrated assumption	Table 4.2 (Loading policy)
Ranging time at dumping for haul trucks (s)	operational confirmation	Table 4.2 (Ranging)
Ranging time at loading for loaders (s)	operational confirmation	Table 4.2 (Ranging)
Additional excavator delay per full load cycle (s)	calibrated assumption	Table 4.2 (Loader variability)
Hauler bay requirement at crushers (bays per hauler)	operational confirmation	Table 4.2 (Crushers and dumping facilities)
Input robustness		
Tail trimming bounds for empirical sampling inputs (percentiles)	model design choice	Vehicle process input construction (Section 4.2.2)

4.3.3. Dynamic Dispatch

The transition from a fixed assignment policy to a DD approach allows haulers to select among the currently operational loaders after each dump. The dispatch decision aims to improve system performance while respecting bench specific targets. Bench depletion is not modeled, therefore the policy is formulated to avoid behavior in which a single attractive location is selected persistently via the bench specific targets.

To balance target compliance with total production, the DD policy is tuned using Optuna, which minimizes/maximizes a single scalar objective function. Therefore both target compliance and total production have been scored on the same variable. Because total production is measured in tons while target compliance is expressed in percentage-point deviation, a penalty weight is introduced to place both terms on a comparable scale. This penalty weight is set to 500 to balance the unit difference in the objective function. Under this scaling, an additional 500 t of production offsets one percentage-point unit of compliance deviation in the objective. As shown in Equation 4.5, the objective function needs to be minimized. Total production has a negative sign, while compliance deviation has as a positive penalty term, so that higher production and lower deviation both reduce the objective value. Compliance de-

viation is evaluated as the sum of the absolute percentage-point deviations between the realized and target bench fractions within a shift.

$$\text{Minimize}\left(\mathbf{F}_{\text{Objective}}\right) = \text{Minimize}\left(-Tonnage_{\text{Hauler}} + \mathbf{W}_{\text{Penalty}} \cdot \text{Compliance Deviation}\right) \tag{4.5}$$

Where:

- $Tonnage_{\text{Hauler}}$: Total tonnage hauled over the evaluation horizon (t).
- $\mathbf{W}_{\text{Penalty}}$: Penalty weight used to scale compliance deviation relative to tonnage (t per percentage-point unit).
- Compliance Deviation: Sum of the absolute percentage-point deviations between realized and target bench fractions within a shift.

Because loading locations are associated with different material characteristics, adherence to the bench fractions serves as a proxy for maintaining the planned material blend. This is particularly relevant because contamination and mud content differ by bench and both these characteristics are not modeled directly in the DES. By evaluating compliance on a fraction basis rather than in absolute tonnage, the dispatch objective assesses whether the relative contribution of the active benches remains close to the planned blend, independent of the total production volume. A limitation of this proxy is that constraints on contaminants such as SiO_2 and mud content are still enforced indirectly through target adherence rather than through explicit quality measurements. As a result, the model may still permit short-term quality violations, for example when several loads from dirty benches arrive in close succession, even if the cumulative bench fractions eventually converge to their targets.

The optimizer runs multiple runs within the DES to optimize the weights of a heuristic formula. The search was performed with the Optuna default Tree Based Parzen Estimator (TPE) search method. This is a search method that evaluates the predefined solution efficiently by incorporating previous evaluations. TPE does this by constructing two probability density functions, l(x) for the good results and g(x) for the bad results (Bergstra et al., 2011). New trials then choose configurations by maximizing the ratio between these two functions. Consequently the search is guided towards regions that have yielded high performance previously, and avoiding ones with low performance. The balance between exploration and exploitation has been achieved via the default hyperparameter values of Optuna. This meant the first ten trials were completely random. After this first phase the splitting threshold (γ) and kernel bandwidth parameters managed the exploration-exploitation balance. These values have been validated in the research of Akiba, the designer of Optuna (Akiba et al., 2019).

The optimization itself consisted of 500 trials with five seeds each, resulting in 2,500 runs in total, with the objective to find the optimal weights for the heuristic formula. The seeds were kept the same for the trials via the CRN technique. To further reduce stochastic noise, total tonnage and compliance deviation were first averaged separately over the five runs of each trial. The final objective value of Equation 4.5 was then computed from these averaged components. This approach reduces the influence of single anomalous runs on the final policy ranking. The variables used have been constructed by analyzing what type of parameters are important when choosing a new loader. One hard rule has been introduced as well. Haulers always choose a far away loader when both crushers are down, as the crush rate is zero at these times. When one crusher is down a similar but less restrictive rule is applied. Haulers are pulled heavily towards far away locations, but it is still possible to travel to near locations when the other parameters of the formula indicate that this is the best choice at that moment. The difference between far away and near location has been made in such a way that every shift has two far away locations and one near location, except for one shift where this split was impossible. This way the rule does not force all of the haulers to one loader when the crushers are both down. This method has been discussed in the mine at a meeting, as this is the best policy during these times but is impossible in the fixed assignment setup. This way the short cycles with high production rates can be prioritized when the load can actually be dumped at the crushers. Formula 4.6 shows the different variables the dispatch formula uses.

$$\text{Score}_{\text{loader}} = w_{\text{travel}} \cdot T_{\text{travel}} - w_{\text{starve}} \cdot T_{\text{starve}} + w_{\text{wait}} \cdot T_{\text{wait}} + w_{\text{target}} \cdot P_{\text{target}} + P_{\text{downtime}} \tag{4.6}$$

Where:

- $\text{Score}_{\text{loader}}$: Cost score assigned to a candidate loader (lower is preferred).
- T_{travel}: Shortest-path travel-time estimate from the truck's current location to the candidate loader (min), computed from the road-network travel-time model.
- T_{starve}: Loader starvation time (min), defined as the elapsed time since the candidate loader last serviced a haul truck.
- T_{wait}: Queueing delay proxy (s), defined as the expected time until the candidate loader can start servicing the arriving haul truck, computed from the current queue and en-route assignments.
- P_{target}: Target signal, expressed in percentage points and computed as the difference between the realized and target bench fractions within the current shift. Loaders at underproducing benches become more attractive, while loaders at overproducing benches become less attractive, with the strength of the signal increasing with the magnitude of the deviation.
- P_{downtime}: Downtime penalty, defined from the current crusher availability state and used to discourage near dispatch decisions during crusher downtime
- $w_{\text{travel}}, w_{\text{starve}}, w_{\text{wait}}, w_{\text{target}}$: Tuned weights (dimensionless).

The weights are tuned using the optimization procedure described above and are applied to Equation 4.6 during simulation. At each dispatch decision, the score is evaluated for every operational loader and the loader with the minimum score is selected. The travel-time term T_{travel} is computed as a shortest-path estimate on the road network. The starvation term T_{starve} is computed from the elapsed time since the candidate loader last serviced a haul truck. The queueing term T_{wait} is computed as an expected delay until service start by combining the current queue and the expected remaining service time of trucks already committed to that loader. The target term P_{target} is implemented as a continuous target signal rather than as a penalty-only term. When a bench is ahead of its target fraction, the signal becomes positive and makes the corresponding loader less attractive by increasing its score. When a bench is behind its target fraction, the signal becomes negative and makes that loader more attractive. The strength of this signal depends on the size of the deviation, so locations that are further behind their targets attract haulers more strongly than locations that are slightly underproducing. This allows the dispatch policy to actively correct deviations in the planned extraction mix during the shift. The downtime term P_{downtime} incorporates crusher availability to discourage dispatch decisions that would increase waiting during crusher downtime. Tuned weight values and performance comparisons are reported in Chapter 5. The optimal set of weights ($\mathbf{w}$) resulting from the TPE optimization process is applied within this four-factor linear cost heuristic. The system calculates this score for every active loader, and the policy selects the loader that yields the lowest overall cost.

4.3.4. Autonomous Vehicles

Every autonomous vehicle experiment is implemented using a consistent modeling approach. Variability in activity durations and payload variables is reduced by squeezing the relevant empirical distributions, reflecting the higher operational consistency attributed to autonomous systems. The squeezing of the distribution is performed with a linear transformation. This transformation preserves the empirical mean of the historical data while forcing the sampled value closer to that mean by dividing the original deviation by a defined Variance Factor (k).

$$P_{sample} = \mu_{emp} + \frac{P_{empirical} - \mu_{emp}}{k} \tag{4.7}$$

Where:

- P_{sample}: The transformed parameter value used in the autonomous simulation.
- $P_{empirical}$: The stochastic value drawn from the historical (manual) distribution.
- μ_{emp}: The mean of the empirical distribution.
- k: The Variance Factor, set to $\mathbf{4.0}$ for these experiments.

This model ensures that the autonomous agents operate with four times the consistency of their manual counterparts (k is set to 4.0 for these experiments), thereby isolating and quantifying the precise value of reduced variance on system performance. The value of 4.0 for the variance squeeze has been adopted from the work of Parreira, as no data on autonomous vehicles was found due to the previously described proprietary nature of this data (Parreira, 2013). A second advantage of autonomous vehicles is their increased utilization through the removal of breaks and shift changes. These gains differ significantly per vehicle type in a setup with one autonomous vehicle. These differences are explained in the next sections per vehicle type.

Autonomous Hauler

This scenario isolates and quantifies the operational impact of replacing a single human operator with an autonomous system. Hauler 36 was selected for this experiment as it was the most utilized vehicle in the historical schedule, being operational every shift in Week 2 and missing only a single shift in Week 1.

The hauler retains its original fixed operational schedule and receives no total operational time increase. This is essential, as the autonomous hauler remains dependent on the schedule and mandatory breaks of the human-operated loaders, whose windows ultimately define the maximum productive time the hauler can achieve. The focus of the experiment is therefore concentrated purely on speed variance, travel time variance, dump time variance and a subsequently explained more efficient break/shift change procedure.

The autonomous hauler is modeled to manage non-productive downtime (breaks and shift changes) more efficiently than its human counterparts. It adheres to the required break/downtime windows, but, unlike the manual haulers, it eliminates the need to physically travel to and from the break room. Instead, the autonomous hauler agent performs a look-ahead check of one hour into the schedule to determine its next load location. It then parks directly at the loader it needs to be at after the break/shift change. This ensures the hauler is correctly positioned at the moment the break/shift change ends. This elimination of non-productive travel is expected to yield a measurable reduction in the restart penalty, a key performance indicator related to the efficiency of resuming operations after mandatory downtime.

Finally, the autonomous driving profile is fully deterministic as the behavioral model is simplified by fixing the Aggressiveness Factor ($\mathbf{AF}$) at $\mathbf{0.0}$. This eliminates the stochastic drift (Δaf) inherent in the manual fleet, ensuring the hauler maintains a perfectly consistent speed profile derived solely from the deterministic road physics ($\mathbf{v}_{\text{pit, state}}$) and safety caps. This isolates the performance improvement resulting from the elimination of human behavior and variability.

Autonomous Loader

This scenario tests the impact of introducing an autonomous loading unit to the operation. Loader 54 was selected as the autonomous agent because of its consistent operational presence in the schedule, being functional every shift in Week 2 and missing only two shifts in Week 1.

The experiment's design is focused purely on consistency, as the autonomous loader does not receive expanded operating hours due to its dependence on operational haulers. The loader's productivity is therefore capped by the performance and breaks of the manual haulage fleet, isolating the value of less variable load cycle time and payload per bucket over increased throughput.

The loader's process variance is reduced by applying the linear variance transformation ($\kappa = \mathbf{4.0}$) to its core activity distributions: the Bucket Cycle Time and the resulting Bucket Payload. This factor ensures the stochastic inputs demonstrate high repeatability, simulating a robotic loading system capable of superior speed and payload consistency compared to a human operator, thereby measuring how this increased predictability affects the queuing and cycle times of the haulers.

Load and Carry

This final experiment evaluates the combined operational benefits of both enhanced precision and increased utilization within the DES of the operation. The only loader assigned LC tasks during the two modeled weeks is LC 53. Therefore LC 53 is configured as an autonomous agent and consequently receives two modeling changes.

The first change is the isolation of the agent's performance variance. The LC's core stochastic activities are modeled with four times the consistency of the manual baseline. The variance factor (k=4.0) is applied to its empirical distributions, including dig time, dump time, travel times (empty and loaded) and bucket payload.

The second change increases the utilization of the LC. The autonomous nature of the LC allows it to operate during periods that were previously blocked by mandatory breaks and shift changes. In addition, periods in which the unit is parked in the baseline schedule are reallocated to productive operation. These changes expanded the operational windows of LC 53 significantly, in contrast to the other autonomous vehicle experiments.

The manual LC operates on an offset schedule designed to maintain continuous crusher feed during standard fleet downtime. During the breaks and shift changes of the standard fleet the feed rate of the LC peaks. As there is no hauler traffic and buffer levels are low, the LC encounters no queuing delays and can dump without interruption. Consequently, the manual baseline already captures the highest-efficiency periods of the shift. Simply extrapolating the average feed rate over additional hours would therefore yield an inaccurate projection. By adding the parked time to the operational time of the LC and removing the breaks and shift changes of this vehicle the autonomous LC worked 13.5 hours more compared to the baseline scenario for both weeks.

4.3.5. Additional experiments

The core scenarios evaluate dispatching and autonomy under the operational constraints present in the Flandersbach mine. In addition, a small set of additional experiments, which loosen one of these constraints per experiment, are evaluated. These experiments provide insight into the upside and trade offs of dispatching and autonomy under relaxed conditions, but they are not used to answer the main research questions because they do not represent the full problem definition.

WT sweep for dispatch trade off behavior

Dispatching in this thesis balances two competing objectives. It seeks to reduce short term inefficiencies by making state dependent assignment decisions, while also maintaining adherence to the shift level bench targets that represent the planned extraction schedule. The relative importance of target adherence is controlled through a target weight parameter, denoted as WT. A higher WT increases the penalty for deviating from targets and therefore prioritizes compliance, whereas a lower WT reduces this penalty and therefore allows the policy to focus more on throughput.

To quantify the throughput compliance trade off, an additional WT sweep experiment is performed. Multiple dispatch runs are executed, each with a different WT value. For each WT setting, the main outputs are the mean weekly total moved tonnage and the mean absolute blend deviation (pp), while the global deviation metric, outlined in the subsequent section, is retained as an additional tonnage-based compliance indicator. This experiment shows how compliance and throughput change across the range of WT values, and therefore quantifies the impact of this parameter.

Resilience under reduced haulage capacity

Dynamic dispatch is expected to provide operational flexibility when the system is disrupted. To test this property, a resilience experiment is performed in which one hauler is removed from the fleet for the full simulated week. The same disruption is applied to both the fixed assignment configuration and the dispatch configuration. This experiment evaluates whether dispatching can maintain better adherence to the planned extraction schedule under hauler failure, and whether any throughput loss can be mitigated through reallocation. The primary reported outputs are total moved tonnage, the fraction-based blend deviation metric, and the tonnage-based global deviation metric. This experiment serves as a proxy for the influence of break downs and the response to these breakdowns for both the fixed assignment and DD strategy.

4.3.6. Scaled autonomous haulage

The autonomous hauler experiment in the core scenario set replaces a single manual hauler with an autonomous hauler. In a mixed fleet, the impact of a low variability autonomous unit can be limited by interaction with manual trucks that have higher behavioral variability. To test whether autonomy effects become stronger when more than one unit is autonomous, an additional scaling experiment

is performed in which three haulers operate with the autonomous driving profile. The three haulers chosen for this experiment are assigned to the same loader on most shifts, to limit their interaction with manual trucks. The purpose of this experiment is to test whether a larger autonomous hauler fleet profits more from the decreased variability due to less frequent interactions with manual haulers.

4.4. KPIs and result aggregation

This thesis evaluates scenarios using a compact set of production, waiting, and compliance KPIs. The KPIs are computed from the DES output logs and are used to compare scenarios on two dimensions that matter operationally at Flandersbach.

First, the KPIs quantify overall material movement and how that movement is distributed across the main material streams. This is required because different technologies can change not only how much material is moved, but also which subsystem produces that tonnage.

Second, the KPIs quantify congestion and operational friction. The mine system is strongly interaction driven, meaning that increased waiting at loaders or along haul roads can remove any potential benefit from faster or more consistent vehicles.

Third, the KPIs quantify adherence to the planned extraction schedule. This is especially relevant for DD, where throughput and schedule adherence can be competing objectives.

4.4.1. Production KPIs

The DES records dump events with a payload quantity (t), a vehicle class (hauler fleet or LC), and a destination category (crusher or stockpile). Production KPIs are computed as sums of dump payloads.

- **Total moved (t).** The sum of all dumped material that goes either to a crusher or to the stockpile. This is the primary throughput KPI. It directly captures the overall material moved by the combined system.
- **Total crushed (t).** The sum of all dumped material delivered to the crushers. This KPI isolates crusher feeding from diversion behavior and is included because the crushers are the downstream constraint that determines how much material can be processed within the simulated horizon.
- **Hauler to crusher (t).** The sum of dump payloads delivered to the crushers by the hauler fleet. This KPI is used to quantify the contribution of the haulage system to crusher feeding.
- **Hauler to stockpile (t).** The sum of dump payloads delivered to the stockpile by the hauler fleet. This KPI is used to quantify diversion behavior and to reveal whether changes in the system cause a larger fraction of haulage output to be routed away from the crushers.
- **LC to crusher (t).** The sum of dump payloads delivered to the crushers by the LC. This KPI is included because LC can change crusher utilization and can displace hauler dumping at the crushers, which can shift work between subsystems rather than increase total throughput.

These stream totals are reported alongside Total moved so that scenario differences can be interpreted as either an overall uplift in moved material or a redistribution between hauler and LC contributions.

4.4.2. Waiting and interaction KPIs

Waiting KPIs capture congestion mechanisms that influence throughput. They are included because scenarios can increase production only when the system can absorb the change without creating additional waiting at loaders, dumping locations, or along haul roads. All waiting KPIs are reported in seconds.

- **LQ pos (s).** The average duration of hauler queuing at loaders, computed using only positive queue events. This metric describes the severity of queuing when a queue occurs.
- **LQ incl0 (s).** The average loader queue time per loading visit, including visits with zero waiting. This metric captures typical waiting per visit and is sensitive to how often queuing occurs.
- **SQ pos (s).** The average duration of stockpile queuing, computed using only positive queue events. This metric describes the severity of stockpile waiting when queuing occurs.

- **SQ incl0 (s).** The average stockpile queue time per stockpile dumping attempt, including attempts with zero waiting. This metric captures typical waiting associated with stockpile diversion and is sensitive to how often stockpile queuing occurs.
- **CQ-LC incl0 (s).** The average queue time of the load-and-carry unit at the crushers per crusher-dumping attempt, including attempts with zero waiting. This metric captures the typical waiting faced by the LC when feeding the crushers and is included because crusher congestion can alter the interaction between LC operation and hauler dumping.
- **Headway (s).** The average positive headway waiting time due to the safety headway logic on road segments and access constraints such as the tunnel. This metric captures travel interaction effects.

Together, these waiting indicators separate production changes driven by improved service rates from production changes that are limited by increased congestion at loaders, dumping facilities, and along the haul roads.

4.4.3. Restart penalty KPI

The restart KPI captures whether queueing is concentrated immediately after scheduled restarts. This metric is included because the mine has regular breaks and shift boundaries, and some scenario changes can alter how vehicles re-enter the system after downtime.

Restart windows are defined as the first 45 minutes after each daily operational start time. The KPI is computed as the ratio of mean loader queue duration in restart windows to the mean loader queue duration outside restart windows. A value above 1 indicates that queueing is higher during restart windows than during the remainder of the operation.

4.4.4. Dynamic dispatch compliance KPIs

DD is evaluated on adherence to shift level bench targets. This KPI is included because DD is designed to allocate limited haulage capacity across competing benches, and the relevant operational objective is not only production volume but also how closely the realized extraction follows the planned schedule. This compliance is evaluated using both a fraction based blend metric and a tonnage based adherence metric.

For each shift, realized hauler tonnage is aggregated by bench. Realized bench totals are converted to realized bench fractions by dividing by the total hauler tonnage in the shift. Shift targets are provided as target fractions per bench. The fraction-based deviations are expressed in percentage points and are used to calculate the mean absolute blend deviation (pp), defined as the mean of the absolute percentage-point differences between realized and target bench fractions, averaged over all active bench-shift combinations. In addition to this fraction-based blend metric, tonnage-based adherence metrics are also reported by aggregating absolute tonnage deviations across benches and shifts.

- **Mean absolute blend deviation (pp).** The mean of the absolute percentage-point differences between realized and target bench fractions, averaged over all active bench-shift combinations.
- **Total target (t).** The sum of planned target tonnes across all active benches and shifts.
- **Total abs deviation (t).** The sum of absolute tonnage deviations across benches and shifts.
- **Global deviation (%).** The total absolute deviation divided by the total target, expressed as a percentage.

4.4.5. Aggregation across runs and weeks

Production KPIs are computed per shift and summed to obtain weekly totals. Scenario outcomes are reported as the mean over ten independent runs. Waiting KPIs, the restart penalty, and the fraction-based compliance KPI mean absolute blend deviation (pp) are computed per run and averaged across runs. Tonnage-based compliance KPIs are aggregated over benches and shifts within each run and then averaged across runs for reporting. This aggregation supports fair comparison between scenarios by reducing sensitivity to stochastic variation in any single run.

4.4.6. Deepdive mode for speed and cycle-time validation

Several validation KPIs in this thesis require a high-fidelity reconstruction of realized haul-cycle speeds and cycle-time structure, aligned at bench level between simulation and historical telemetry. To support this, the simulator includes an optional `deepdive` mode, implemented as a dedicated post-processing routine that performs a full dataset versus simulation comparison using the detailed event logs produced by the DES. The deepdive routine is computationally expensive and therefore not executed for every experiment run, but it is used for the specific calibration and validation comparisons that underpin the reported speed and cycle-time diagnostics.

Operationally, deepdive takes as inputs: the simulation `movement_log` and `edge_speed_log` (stamped vehicle movement and segment-level information), the historical cycle dataset containing cycle times and derived loaded/empty speeds, and the directed road network graph G used for routing. Deepdive then reconstructs comparable cycle representations per bench label by applying the same bench naming logic to both datasets, and aggregates results by location for reporting.

Because the historical speed values are derived from recorded distance and travel time and may include corrupted cycles where travel time approaches zero, deepdive applies a corruption filter that removes historical cycles with unrealistic speeds ($V_{\text{loaded}} > 40$ km/h or $V_{\text{empty}} > 40$ km/h). In addition, the deepdive comparison applies a simulation-side outlier cutoff by excluding reconstructed simulated cycles exceeding 60 minutes. Summary statistics are then computed per location, including mean, median, and percentile bands (P10/P90), and exported to the thesis KPI workbook.

The speed-validation tables and the supporting cycle-time diagnostics reported in the Results chapter (and Appendix I) are produced from this deepdive procedure, rather than from the standard fast KPI pipeline. This design ensures that the reported speed comparisons reflect the most detailed and internally consistent reconstruction available from the simulator, at the cost of runtime.

4.5. DES Visualization

To support implementation verification of the constraints listed in Table 4.2 and to aid debugging of scenario logic, a custom playback visualization was developed using the matplotlib.animation library.

To preserve the computational speed of the model the visualization reads the pre-generated event logs and trace data generated during the simulation. From these logs and data the visualization reconstructs: the movement of the haulers, the workings of the LC, the status of the crushers and the status at the loaders. A screenshot of this visualization is depicted in Figure 4.1. A zoomed-in version of the status panels (crusher state and buffers, bench queue status, and break-room occupancy) is provided in the Appendix G for readability. The different colors and panels represent:

1. **Vehicle States:** Agents are color-coded to visualize their current status: **Blue** (Empty travel), **Orange** (Loaded travel), **Red** (Queuing), and **Magenta** (Processing/Loading/Dumping). This allows for immediate identification of flow interruptions.
2. **Buffer Dynamics:** Panels monitor the Crusher Buffer levels, verifying that the vehicles divert to the Stockpile when the crushers buffers reach their thresholds (indicated by the HUD turning orange), and shows a list of haulers in the break room.
3. **Queue Formation:** A specific ”Bench Status” panel monitors the build-up of queues at every load location, it displays the quantity of haulers currently being loaded/in the queue per bench.

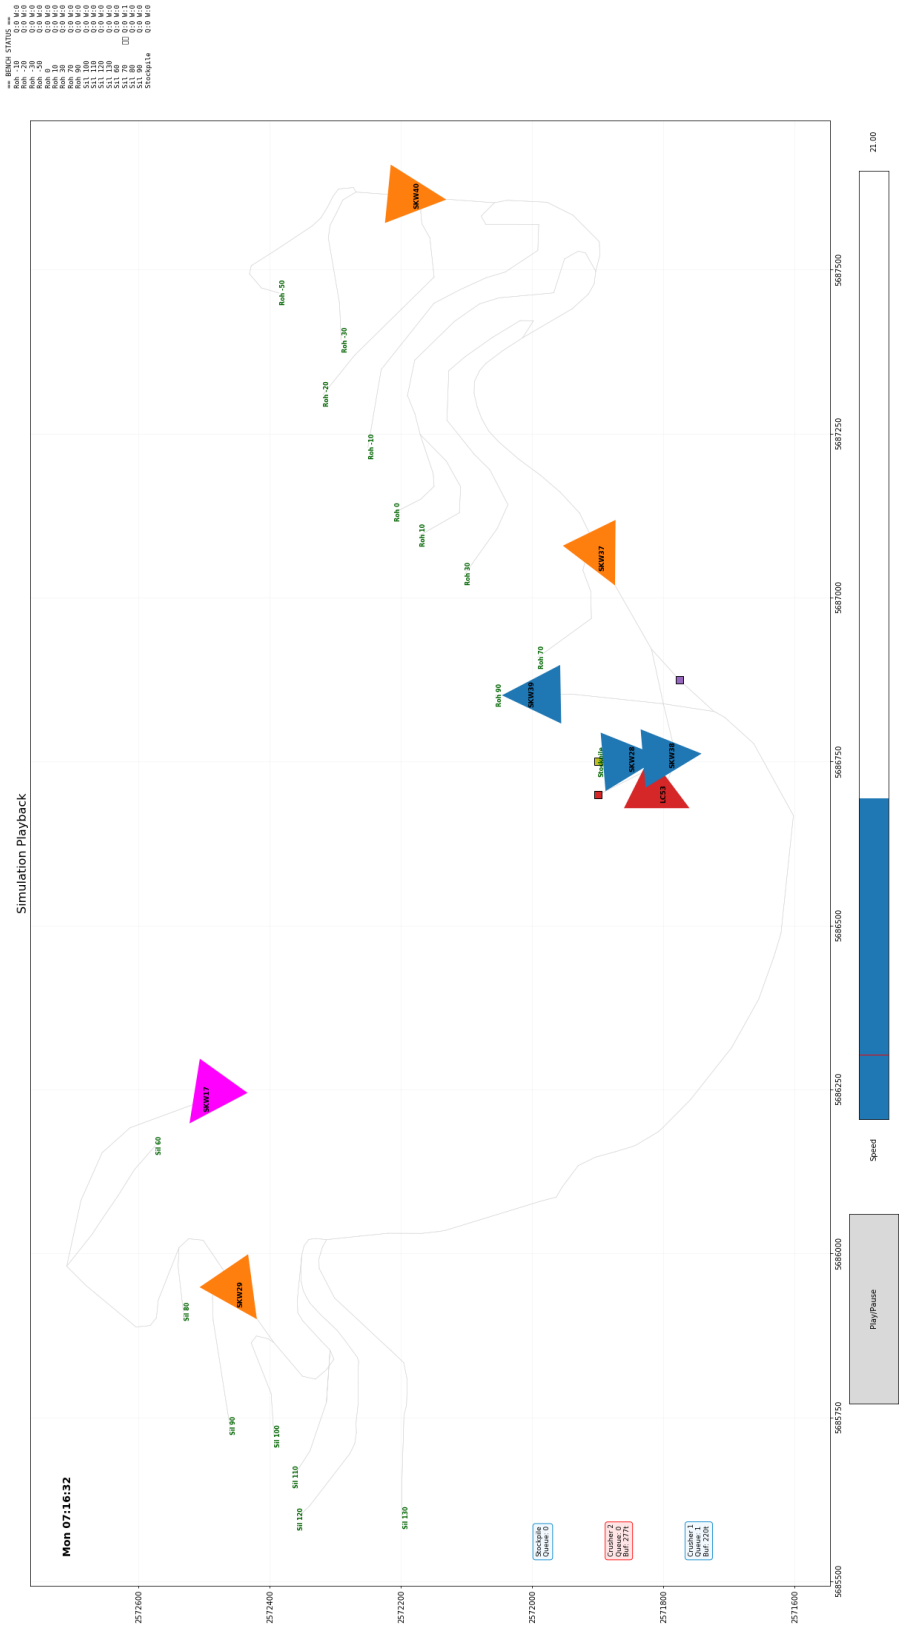

Figure 4.1: Playback visualization of the DES environment. The interface shows the 2D road network with color-coded vehicle states and panels monitoring crusher buffer levels and bench queue status.

5

Experiments and Results

5.1. Chapter overview and evaluation framing

This chapter contains the validation and experimental outcomes of the DES model of the Flandersbach mine used to answer the research questions of this thesis. This chapter consists of three different result layers.

First, the chapter summarizes the data analysis output used to parameterize the DES. This includes the construction of empirical distributions derived from telemetry data and production logs for haulers, loaders and LC, as well as the pooling and trimming decisions used to obtain representative distributions. Detailed diagnostics and distribution plots are provided in the corresponding appendices and are referenced to support modeling choices.

Next, the road network representation of the mine used in the DES is presented. The network provides an overview of the mapping and connectivity between operational locations such as the benches, crushers, the stockpile and breakroom/petrol station.

Validation is reported using a calibration and generalizability framing. Week 2 is treated as the calibration period and was used to align the model with observed performance. Week 1 is then used as a check to evaluate whether the calibrated model remains representative when applied to a different week. Validation is presented from component to system level, starting with input plausibility checks against telemetry derived distributions and progressing to local operational behaviour and overall production outcomes.

Finally, the chapter reports the scenario experiment results for the dispatching and autonomy configurations defined in Chapter 4. This chapter focuses on outcomes and comparisons, while the formal scenario definitions, KPI calculations, and the experiment descriptions are provided in the Methodology chapter to avoid duplication. For the experiments production performance is measured through total moved tonnage and moved tonnage per material flow. Operational stability is reported through target adherence, queue formation, road interaction waiting and restart behavior. This distinction is used to keep the production and operational stability effects seperate.

As a companion to the architectural overview of the DES in Figure E.1, Appendix H provides a second schematic focused on the reporting structure of the results chapter (Figure H.1). Together, these figures distinguish where the main submodels act in the model from how the model is parameterized and which outputs are reported in Chapter 5. The results schematic is included as a reading aid to clarify the hierarchy of the chapter rather than as a separate result.

5.2. Inputs: empirical basis and final parameterization

This section reports the data-analysis outcomes that provide the input parameters of the DES. The empirical distributions used in the DES are obtained by applying the preprocessing steps described in Chapter 4. The focus in this section is on the resulting datasets and their implications for parameteri-

zation. Where measurements are incomplete or inconsistent across the fleet, the resulting choices are stated explicitly, including pooling, substitution, and the use of fixed values.

The most important modeling decisions regarding telemetry data are that hauler payload emerges from loading logic rather than from a direct hauler payload distribution. Loader behavior relies on derived profiles, with an additional delay in one case. LC behavior is represented through crusher feeding cycles only. These points are relevant when interpreting the validation and scenario results reported later in this chapter.

To make the empirical basis transparent before discussing each subsystem, Table 5.1 summarizes the final sample sizes and the role of each dataset in the study. It distinguishes between streams used directly as stochastic model inputs and streams used for validation of the DES. The remainder of the section then documents the resulting input basis per asset class, and references the corresponding appendices for the supporting evidence regarding the modeling choices.

Table 5.1: Sample sizes for DES inputs and key validation references.

Asset / stream	Data basis	Sample size	Used for
Haulers	Bench-matched cycles with speed information	19,901 cycles	Validation of hauler travel speeds by bench (reporting/validation reference; not used as a direct DES input).
	Dump time (CAT 777D only)	6,121 records	DES dump-time sampler for haulers. Source: #16 CAT 777D (2,422) and #17 CAT 777D (3,699).
	Payload reference (variable payload only)	32,305 records	Validation reference for simulated hauler payloads generated by bucket-by-bucket loading logic (not a direct hauler input). Source: #28 (6,804), #29 (6,402), #32 (3,667), #36 (7,435), #37 (7,997).
Loaders	Bucket cycle time input (sequence 2–5)	100,097 bucket events	DES loader bucket cycle-time sampler during hauler loading (sequence 1 excluded due to first-bucket effect). Source: LG54 (45,119) and LG55 (54,978).
	Bucket payload input (sequence 1–5)	126,699 bucket events	DES loader bucket-payload sampler during hauler loading (restricted to full-bucket portion of the sequence). Source: LG54 (57,097) and LG55 (69,602).
	Full-load cycle comparison basis	LG14: 12,546; LG53: 8,099 cycles	Reporting/validation comparison of full-load cycle times across equipment types, supporting the excavator penalty setting.
	Loader profiles derived from LG54/LG55	–	LG53 uses the LG54 bucket cycle-time and bucket payload distributions. LG14 uses the LG54 distributions and applies an additional 20 s delay per full load cycle.
Load-and-carry	Crusher-feeding stream (Interim A $\rightarrow$ Crusher 185186)	8,333 cycles	DES LC samplers for payload, dig time, dump time, loaded travel time, and empty travel time, pooled across LC54 and LC55. Source: LC54 (4,542) and LC55 (3,791).

5.2.1. Haulers: telemetry availability, pooling decisions, and final input distributions

The telemetry data of VisionLink did not record speed data directly, therefore the speed data is derived from travel time and travel distance for loaded and empty states. After addressing missing distances and assigning benches to the full dataset a total of 19,901 rows remained in the dataset. This dataset is used later as a validation method to assess whether the hauler movement in the model is conform reality.

The haulers in the DES use dump time as their sole empirical input. Dump time recordings are not available for every hauler and are not consistent while the manufacturer specifications indicate otherwise,

as can be seen from Table 3.1. The Komatsu haulers exhibit dump time values that reflect time spent at the crusher, rather than the tipping event. The Caterpillar 777D haulers (Hauler 16 and 17) were the only haulers with dump times in line with the manufacturer specifications, which indicate limited differences between all haulers. Therefore the telemetry data of these haulers is used by every hauler in the DES. This yields 6,121 dump-time observations used for the dump-time sampler (2,422 records for Hauler 16 and 3,699 records for Hauler 17). The resulting input distribution is provided in Appendix A.

Payload is not a direct input of the DES, but is used to validate the loading process of the DES and the associated payload threshold. Payload records exist for every hauler, however several vehicles show unrealistic constant payload values. For this purpose, the payload reference set is constructed using the haulers with a variable payload. This includes Hauler 28, 29, 32, 36, and 37, providing 32,305 payload observations. The empirical payload basis and the corresponding comparison figures are provided in the Hauler Appendix A. Thus the hauler payload output of the DES is not imposed directly, but emerges from the interaction between loader bucket payloads and the loading threshold.

5.2.2. Loaders: bucket sequence filtering, cycle measurements, and final input distributions

Within the DES haulers are loaded bucket by bucket, using empirical inputs for both bucket cycle time and payload. Data on bucket level is only available for loader 54 and 55. Therefore these loaders form the empirical basis for the other two loaders as well, through derived profiles described below.

Two sampling windows are used to derive the distributions used in the DES. Bucket cycle times are sampled using bucket events with sequence numbers 2-5, yielding 45,119 events for Loader 54 and 54,978 events for Loader 55. Buckets with sequence number 1 have a significantly longer cycle time due to the ranging time of the hauler. Bucket payloads are sampled using events with sequence numbers 1–5, yielding 57,097 events for Loader 54 and 69,602 events for Loader 55. This window captures the full-bucket portion of the loading sequence and excludes rare high-sequence events that are typically associated with partially filled buckets. The evidence supporting these sequence restrictions, and the resulting input-versus-simulation comparisons, are provided in the Appendix B.

Loader 53 and loader 14 did not contain the bucket-level payload and bucket cycle-time information required for the bucket-by-bucket loading model. Therefore, their loading behavior is parameterized using the empirical distributions of the other loaders. Loader 53 is the same wheel-loader model as loader 54, and is therefore modeled directly using the loader 54 bucket payload and bucket cycle-time distributions. Loader 14 is an excavator and exhibits longer full-load cycle times than the wheel-loaders. To reflect this difference while retaining the same bucket-level sampling structure, loader 14 is modeled using the loader 54 bucket-level distributions with an additional delay per full load cycle equal to the mean difference observed between the full-load cycle time distributions of loader 14 and loader 53 (20 seconds). The full-load cycle comparison used to support these substitution comprises 12,546 usable cycles for loader 14 and 8,099 usable cycles for loader 53, the corresponding figures are provided in Appendix B. These substitution mean that the loader profiles for loader 14 and loader 53 represent the best available approximations for the simulated weeks, but are not exact reconstructions.

5.2.3. Load-and-carry: proxy variables, other activities, and final input distributions

The LC vehicle combines loading and hauling, consequently it requires multiple inputs in the DES, including payload, dig time, dump time and travel time in loaded and empty state. The wheel-loader employed as LC in the simulated weeks does not provide a complete activity level dataset in the VisionLink data. Therefore, LC input distributions are derived by combining the activity level datasets of LC 55 and LC 54, with analogous manufacturer specifications as can be seen from Table 3.2. Supporting figures and comparisons are provided in Appendix C.

To ensure that the LC input basis represents crusher-feeding operation, the dataset is restricted to cycles corresponding to the crusher feeding activity. This yields 4,542 cycles for LC54 and 3,791 cycles for LC55, resulting in a pooled basis of 8,333 cycles used to parameterize LC payload and activity-time samplers in the DES.

A key consideration for LC behavior is that LC vehicles have a dual role in practice. Feeding the crusher is the primary task, but stockpile management and main dumping area cleaning also consume a substantial and variable fraction of operating time. This appears in the telemetry as periods that do not correspond to the main cycle sequence and are summarized as “other activities”. Other activities are defined as within-shift gaps in the downloaded cycle data that exceed one minute. The distribution of this quantity shows that the operational time spent on non-crusher-feeding tasks is highly variable. Due to the lack of data for LC 53 during the two simulated weeks the precise share it spent on other activities was not possible to retriev. The corresponding ”other activities” plot is provided in Appendix C. These ”other activities” were not modeled during the simulated weeks, the implications of this modeling decision are discussed in 6.

5.3. Road network representation and verification

The DES represents the road network of Flandersbach as a bidirectional graph, with the exception of the single lane tunnel, that connects benches to the main dumping area and the petrol station and break room. Figure 5.1shows the 2D overview of this network, including the facilities and benches. This figure confirms that the topology and labeling implemented in the DES match the mine layout, and that each bench has a valid path connecting it to the facilities. Consequently this network validation is treated as a structural check of the DES.

Additionally, the network computes segment lengths and elevation differences, which determine the grade per segment and influence the modeled travel speeds. A 3D representation of the same network is provided in Appendix D (Figure D.1). This view supports qualitative verification that vertical structure is consistent with the pit layout, and that grade effects can be applied coherently across the haul roads.

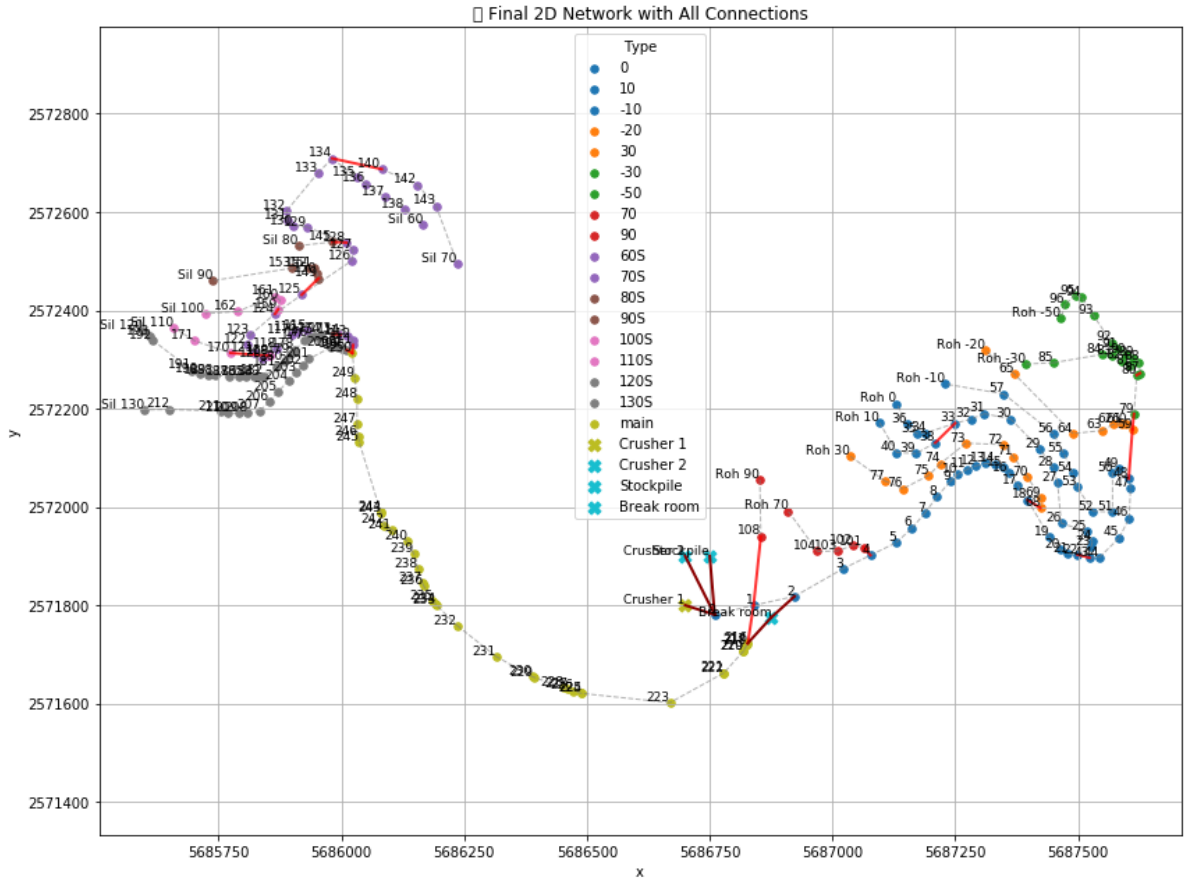

Figure 5.1: Two-dimensional overview of the directed haul-road network used in the DES. Note that the coordinate system of this figure follows the site-provided coordinate reference used to construct the network and therefore differs from the coordinate system of the aerial mine image (Figure 3.1)

5.4. Model validation results

Validation is performed against two operational weeks to ensure the DES provides a robust representation of the Flandersbach mine. Week 2 (24-02-2025 to 28-02-2025) is used as the calibration period, in which model parameters are aligned within plausible ranges to reproduce observed system behavior. Week 1 (17-02-2025 to 21-02-2025) is then used as a generalization test, where the calibrated model is evaluated on a different operating week without further adjustment.

Using multiple weeks is important because mining operations are inherently dynamic and stochastic. Even under similar operating conditions, production outcomes can vary substantially due to random events. Therefore, a deviation in a single shift does not necessarily indicate a model deficiency. Deviations that are persistent in sign or recur systematically across shifts and across weeks are more indicative of structural bias in the model representation.

The validation presented in the subsequent section is presented in hierarchical fashion. First, component level behavior is assessed through input plausibility and operational logic checks, ensuring that empirical sampling and key rules generate realistic model behavior. Next, system level performance is evaluated through production comparisons at week, shift and bench level . Finally, the deepdive checks are used to relate local deviations to underlying cycle-time and speed behavior deviations.

5.4.1. Component level validation

The component level validation focuses on the reproduction of the intended process behavior of the DES compared to reality. First, the empirical input distributions are checked against the simulated draws of the DES. Second, the rule driven outcomes that are not directly sampled, such as hauler payloads and break behavior, are evaluated to verify that the implemented rules produce realistic outcomes. The supporting figures for these checks are provided in Appendix I. The presented figures are from Week 2, equivalent figures for Week 1 are similar and are omitted for brevity.

5.4.2. Input sampling checks

The sole input for the hauler process consists of the empirical dump time distribution. The comparison between the empirical distribution and the simulated draws indicates that the simulated values reproduce the shape and mean of the empirical input, confirming that the stochastic representation preserves variability in the DES (Appendix I, Figure I.1)

For the LC process, the simulator samples payload and digging, travel, and dumping durations from the processed empirical inputs. The comparison between the empirical distributions and the simulated draws indicates that the simulated values reproduce the shape and mean of the empirical inputs across all activities, confirming that the stochastic representation preserves variability in the DES (Appendix I, Figure I.2).

The loader process takes bucket payload and cycle time samples from the empirical distribution. The comparison between the drawn samples and the distribution shows comparable mean and shape and is presented in Appendix I, Figures I.3 and I.4. Additionally, a full load cycle comparison across all modeled loaders is used as a consistency check for the derived loader profiles. The resulting distributions show that loader 14 exhibits longer full-load cycle times than the wheel-loader profiles, in line with the additional delay applied to represent the excavator loading process (Appendix I, Figure I.5).

5.4.3. Rule driven outcome checks

Rule driven outcome checks confirm that implemented rules produces realistic process outcomes. This is particularly relevant for quantities that are not sampled directly but emerge from the interaction between model rules and sampled inputs.

Hauler payload is not sampled directly in the DES. Instead, payload results from the bucket-by-bucket loading process and the loading threshold, which has been set to 87 ton via multiple iterations. A comparison between the simulated payload distribution and the telemetry-based payload reference for haulers with reliable payload variability shows that the model reproduces the typical payload range and central tendency, while small distributional differences remain due to the simplification that partially filled buckets are not represented (Appendix I, Figure I.6). This check confirms that the implemented loading logic yields plausible payload outcomes at the fleet level.

Break and shift-change behavior is represented through explicit schedule logic and early-break rules intended to avoid unrealistically short break-room dwells. A dwell-time check is used to verify that the resulting break-room time allocation is consistent with the planned break structure, including the longer overnight closure period. Appendix I, Figure I.7, shows the resulting break time distributions, for an early break threshold of 14 minutes. After multiple iterations this value was found to approximate the real break times best.

To verify that key operational rules are executed as intended during simulation, qualitative behavior traces are included. These figures illustrate that diversion and buffering logic are followed consistently, including diversion of haul trucks to the stockpile when crusher conditions prevent dumping, and the resulting evolution of stockpile accumulation over time (Appendix I, Figures I.10 and I.11). Activity pie charts summarize how simulated time is allocated across the main task categories for both the hauler fleet and the LC, providing a plausibility check that the dominant time shares are spent in the expected activities (Appendix I, Figures I.8 and I.9). These checks support that the DES reproduces the intended process behavior and rule consequences.

Additional snapshots from the playback visualization are provided in Appendix G and confirm that the implemented rules produce the intended local behavior, including day start conditions, early-break behavior, break-room accumulation during breaks, and shift change load location changes.

5.4.4. System-level production validation

Global Production

This macro level validation compares the total material moved and processed in the DES against the production records of the mine for both simulated weeks. Table 5.2 summarizes the aggregated production metrics for Week 1 (generalization) and Week 2 (calibration). Across both weeks, the model exhibits a consistent positive bias in total moved tonnage (Week 1: +6.0%; Week 2: +6.8%) and total crushed tonnage (Week 1: +10.1%; Week 2: +10.3%). The persistence of this deviation across both weeks suggests a structural mismatch between the simulation and reality.

The contribution to this bias differs between subsystems. For the hauler fleet, the simulated total hauled tonnage exceeds the observed totals by +2.9% in Week 1 and +5.7% in Week 2. This difference is primarily reflected in the hauler-to-crusher stream, where the simulation produces +6.6% and +9.6% more crusher feed than recorded for Week 1 and Week 2, respectively. In contrast, the hauler-to-stockpile stream is underestimated in both weeks (Week 1: −5.9%; Week 2: −6.9%), indicating that the model tends to route additional tonnage to the crushers, which cannibalizes on the hauler-to-stockpile stream.

The LC production numbers show the largest and most consistent deviation, producing +23.2% more tonnage than observed in Week 1 and +30.9% more in Week 2. This is consistent with the modeling scope described in Section 5.2, as the LC is represented as a dedicated crusher-feeding unit, while in practice it also performs auxiliary duties.

To contextualize the weekly totals, shift-level real versus simulation comparisons are provided in Appendix I for both weeks (Tables I.3 and I.4). These tables report real, simulated, and difference values per shift for hauler-to-crusher, hauler-to-stockpile, load-and-carry-to-crusher, total moved, and total crushed, and are included to verify that the weekly deviations are not driven by a single or a subset of outlier shift(s).

While the DES was calibrated using Week 2 data, the error magnitude and direction of the errors remains similar for Week 1. This suggests the model is generalizable to periods outside the calibration window.

Table 5.2: Dual-week production summary: real-world totals versus simulation (mean over runs).

	Week 1				Week 2			
Metric	**Real (t)**	**Sim (t)**	**Delta (t)**	**Dev (%)**	**Real (t)**	**Sim (t)**	**Delta (t)**	**Dev (%)**
Total Moved	145,953	154,769.9	+8,816.9	+6.0%	141,680	151,309.0	+9,629.0	+6.8%
Total Crushed	109,055	120,036.1	+10,981.1	+10.1%	112,802	124,427.0	+11,625.0	+10.3%
Hauler Fleet	123,254	126,802.6	+3,548.6	+2.9%	121,190	128,088.2	+6,898.2	+5.7%
to Crusher	86,356	92,068.9	+5,712.9	+6.6%	92,312	101,205.6	+8,893.6	+9.6%
to Stockpile	36,898	34,733.7	-2,164.3	-5.9%	28,878	26,882.6	-1,995.4	-6.9%
Load & Carry	22,699	27,967.3	+5,268.3	+23.2%	17,743	23,220.8	+5,477.8	+30.9%

Local production

To assess whether the global deviations are uniformly distributed across operational locations, hauler production is compared on bench level per shift. Bench-level deviations for Week 1 and Week 2 are reported in Appendix I (Tables I.1 and I.2). The results show deviations are concentrated in a subset of benches that exhibit repeated, extreme positive deviations. Other benches remain accurate or alternate between over- and underperformance across shifts, indicating no systematic bias at these benches.

In Week 2, the largest systematic contribution is associated with Rohdenhaus bench Roh 90. This bench shows very large positive deviations in multiple early-week shifts, including Monday Early (+57.1%, +2,375 t) and Tuesday Early (+44.4%, +2,021 t), resulting in a large positive bench total (+21.4%, +4,540 t). In the same week, Roh −50 and Sil 60 also exhibit several shifts with substantial positive deviations (e.g., Roh −50 Thursday Early +20.6%, +877 t; Sil 60 Friday Early +35.0%, +624 t). In contrast, Sil 70 remains close to observed production across multiple shifts, typically within a few percent, indicating that not all locations contribute equally to the observed system-level bias.

Week 1 exhibits a different local deviation pattern, underscoring the week-specific nature of operational variability. The strongest repeated deviations occur at Sil 80, including Tuesday Late (+47.5%, +1,126 t) and Wednesday Early (+58.9%, +1,245 t), resulting in a sizeable positive bench total (+9.2%, +2,796 t). Sil 110 shows a single large deviation in the only shift in which it is active (+37.2%, +920 t). At the same time, several locations show mixed performance across shifts, with deviations changing sign over the week.

For the subsequent cycle time and speed analysis, benches are interpreted in three qualitative categories based on Tables I.1 and I.2: (1.) benches with repeated positive deviations that dominate the weekly bias, (2.) benches with sparse activity or large sign changes across shifts, and (3.) benches that remain close to observed production in most active shifts. This grouping is used to evaluate the impact of the speed model on the observed production deviations.

5.4.5. Speed and cycle time

To identify potential drivers of the localized deviations described in the previous subsection, a deepdive analysis is performed on haul-cycle structure and achieved travel speeds. Empty and loaded travel constitute the largest share of hauler operating time in the DES, meaning that optimistic travel speeds can reduce modeled cycle times and translate directly into higher production. Therefore this subsection examines whether local production deviations are consistent with deviations in cycle times and speed behavior, and where this explanation breaks down.

Cycle-time comparisons are summarized in Table 5.3, which contrasts historical baselines with the specific weekly outcomes and the simulated outcomes. Detailed speed validation for loaded and empty travel is provided in Appendix I (Tables I.5 and I.6), and spatial distributions of stamped speeds on the network are shown in Figures I.13 and I.12. The speed validation statistics presented in Appendix I were generated using the Deepdive validation routine described in Section 4.4.6. This routine reconstructs cycle-level speed representations from detailed simulation logs and aligns them with filtered historical telemetry data for consistent comparison. Together, these diagnostics provide a physical basis for interpreting repeated bench-level overproduction as either too many trips (cycle times too low), implausible speed behavior, or interaction effects that are not captured by speed alone.

For some benches, the expected causal chain is observed. An illustrative case is Roh 90 in Week 2, where modeled travel speeds exceed the measured values and the simulated cycle time is lower than observed. This combination is consistent with the large positive production deviation of Roh 90 reported in Table I.2. Similar behavior is observed for Sil 120 in Week 2, where higher modeled speeds correspond to lower simulated cycle times and positive production deviations. In these cases, speed behavior provides a sufficient explanation for the direction of the production error.

However, the deepdive also highlights that speed alone does not explain all deviations. Roh 90 in Week 1 provides a counterexample: despite substantially higher modeled empty speeds relative to the measured Week 1 conditions, the simulated cycle time is not faster, while the location still overproduces. A similar contradiction appears in locations that are otherwise reproduced accurately at production level. For multiple benches, simulated speeds are lower than measured while simulated cycle times remain comparable or even slightly faster. Finally, there are cases where production and cycle-time signals diverge strongly. For example, Sil 110 in Week 1 shows a large positive production deviation while the simulated cycle time is not faster than observed.

Table 5.3: Cycle Time Validation: Historical Baselines vs. Specific Weekly Conditions

Location	Real Data Historical	Real Data Week 1	Real Data Week 2	Simulation Week 1	Simulation Week 2	Deviation
Recurring Locations						
Roh -10	19.9	24.9	25.3	24.6	24.2	-1.2% / -4.3%
Roh -20	20.4	20.2	21.0	22.1	21.8	+9.4% / +3.8%
Roh 90	14.8	10.9	12.5	11.0	11.8	+0.9% / -5.6%
Sil 70	27.3	35.3	28.5	30.0	29.5	-15.0% / +3.5%
Sil 110	18.8	20.3	19.4	21.1	22.2	+3.9% / +14.4%
Week 1 Exclusives						
Sil 80	23.8	24.2	–	24.2	–	+0.1%
Sil 100	21.3	22.1	–	22.5	–	+1.8%
Roh -50	24.0	22.7	–	22.3	–	-1.8%
Week 2 Exclusives						
Sil 120	20.9	–	21.9	–	20.5	-6.4%
Sil 60	27.0	–	25.0	–	27.8	+11.2%

Note: Deviations are calculated against the specific week's real data. Historical data provided for context.

5.5. Sensitivity analysis results

This section reports the outcomes of the one-factor-at-a-time sensitivity analysis defined in Chapter 4 (Table 4.5). Each setting changes a single parameter relative to the Baseline configuration while all other inputs remain unchanged. CRN is used across settings so that differences in results reflect the parameter change rather than stochastic variation. Autonomy and dispatch scenarios, and the additional experiments, are not part of the sensitivity analysis and are reported in the scenario experiment results section.

Sensitivity of production metrics

Production sensitivity is summarized for Week 1 and Week 2 in Appendix I (Tables I.7 and I.8). Following the structure defined in Chapter 4, results are interpreted by mechanism category: crusher parameters (capacity and downtime), transportation parameters (speed and behavioural variability), loading and dumping parameters (ranging and excavator penalty), and input robustness (trimming).

Across both weeks, the largest production changes occur under transportation settings and under large throughput changes. The haul-truck setting (HT_95) increases total moved tonnage in both weeks (+3.33% in Week 1 and +3.14% in Week 2), while (HT_80) yields the largest decrease (-6.29%

and -6.43%, respectively). Changing crusher processing rate shows the same directional consistency: (CR_TPH_1200) increases total moved tonnage (+1.59% and +1.26%) and (CR_TPH_800) decreases it (-1.77% and -1.44%). Changes in buffer and bay configuration ((BUF_200), (BUF_400), (HBAYS_2)) lead to smaller but still consistent shifts.

The speed-related settings also show clear effects on production. The global speed increase setting (SPEED_MULT_1p1) increases total moved tonnage by +3.94% in Week 1 and +3.71% in Week 2, while (SPEED_MULT_0p9) decreases total moved tonnage by -3.80% and -4.06%, respectively. The grade-effect multiplier settings show a similar directional effect, with (SPEED_B_MULT_1) increasing tonnage (+2.40% in Week 1 and +1.87% in Week 2) and (SPEED_B_MULT_5) decreasing it (-4.55% and -3.50%). These settings scale the grade penalty in the speed model, reducing or increasing the effect of slope on loaded vehicles and scaling the empty-vehicle effect accordingly. The variability parameter (SPD_KAPPA_0p05) also increases production in both weeks (+2.08% and +1.82%), while (SPD_KAPPA_0p25) decreases it (-0.77% and -1.73%). These results show that transportation assumptions are among the most influential drivers of total moved tonnage in the model.

Loading and dumping settings, including ranging-time variants ((RANGH_5), (RANGH_15), (RANGL_10), (RANGL_30)) and the excavator delay variants ((LG14EX_0), (LG14EX_40)), generally produce smaller changes in total moved tonnage than the dominant transportation and throughput factors. The loader squeezing variants ((SIG_LG54_2), (SIG_LG54_3), (SIG_LG54_5), (SIG_LG54_6)) also remain relatively close to the Baseline outcomes.

Finally, trimming primarily matters at the extremes. Moderate trimming adjustments ((TRIM_5p0_95), (TRIM_1_99)) produce limited changes, while the no-trim case (TRIM_0_100) produces a much larger change in Week 1 (-3.15%) and a smaller but still negative change in Week 2 (-0.95%). This setting is therefore best interpreted as a robustness stress test rather than a representative operating assumption.

Sensitivity of queues and waiting

Waiting-time sensitivity is reported in Appendix I (Tables I.9 and I.10). These KPIs are used as supporting indicators to contextualize production differences, rather than as calibration targets.

Most settings leave loader queues and headway waiting within a similar range to the Baseline. The clearest changes in queue behaviour occur under transportation variability settings, where behavioural variability affects the interaction between arrivals and service at loaders. This is most visible for (SPD_KAPPA_0p05) and (SPD_KAPPA_0p25). In both weeks, (SPD_KAPPA_0p05) substantially reduces loader queuing, while increasing headway waiting and restart penalty. Conversely, (SPD_KAPPA_0p25) increases loader queuing while reducing headway waiting. Crusher capacity and downtime settings mainly shift waiting at dumping facilities and the stockpile, consistent with changes in diversion behavior and the timing of crusher feeding under altered capacity or availability.

The speed-related settings also affect waiting behavior, but less strongly than the corresponding production changes. Across both weeks, the global speed settings and the grade-effect multiplier settings shift loader queuing, crusher queuing for load-and-carry, and headway waiting, but these effects remain moderate compared with the more extreme queue responses under variability settings and trimming extremes.

A notable robustness case is (TRIM_0_100) in Week 1, which produces extreme stockpile queue values (1397.24 s positive-only and 220.59 s including zeroes). This indicates that including the full empirical tails can introduce rare but dominant events that heavily affect waiting metrics, which supports the use of trimmed empirical inputs in the baseline configuration.

Summary of influential assumptions

Overall, the sensitivity analysis indicates that the dominant drivers of total moved tonnage are transportation settings that influence cycle-time efficiency and major crusher capacity changes that constrain system throughput. In particular, the speed settings, grade-effect settings, behavioural variability settings, and haul-truck payload assumptions can materially change production across both weeks. In contrast, several loading and dumping settings affect the composition of waiting time and local congestion more than total moved tonnage. These findings support the Baseline parameterization used in the

subsequent scenario comparisons and identify which assumptions should be treated as high-impact when interpreting scenario results.

5.6. Core scenario experiment results

Table 5.4 summarizes the core scenario outcomes for both simulated weeks. The table reports total moved tonnage as the primary KPI, accompanied by stream-level totals for hauler to crusher, hauler to stockpile, and LC to crusher. It also includes operational indicators such as loader queue time, headway waiting, and the restart penalty. This structure allows differences between scenarios to be interpreted based on both aggregate production and the distribution of work between the hauler fleet and the LC unit.

Across both weeks, the autonomous LC case yields the largest increase in total moved tonnage. In Week 1, total moved increases from 154,769.9 t in the Baseline to 166,580.7 t, representing a 7.63% gain. In Week 2, total moved increases from 151,309.0 t to 163,811.6 t, an 8.26% improvement. The stream totals indicate that this increase is primarily driven by higher LC crusher feed. In Week 2, LC to crusher increases from 23,220.8 t to 35,404.9 t, while hauler to crusher decreases from 101,205.6 t to 95,412.6 t and hauler to stockpile increases from 26,882.6 t to 32,994.1 t. A similar redistribution is visible in Week 1. These results show that the autonomous LC gain is not a uniform uplift across all streams but rather a shift in how material is moved and directed.

The autonomous hauler case produces a smaller but consistent increase in total moved tonnage in both weeks, with gains of 1.52% in Week 1 and 1.81% in Week 2. The stream breakdown shows that this improvement is expressed mainly through higher hauler crusher feed. In Week 2, hauler to crusher increases from 101,205.6 t in the Baseline to 103,434.9 t, while LC to crusher decreases slightly from 23,220.8 t to 22,904.2 t. The autonomous loader case yields a modest increase in both weeks of 0.72% and 0.64% respectively, with only small changes in the stream totals relative to the Baseline.

Dynamic dispatch reduces total moved tonnage relative to the Baseline in both weeks, showing a 0.51% reduction in Week 1 and a 0.48% reduction in Week 2. The waiting indicators provide context for this reduction. In both weeks, loader queue time increases substantially under dynamic dispatch. In Week 2, for instance, LQ pos increases from 133.46 s to 192.75 s and headway waiting increases from 23.47 s to 28.58 s. At the same time, the restart penalty increases relative to the Baseline in Week 2, from 1.41 to 1.74. This shows that the small reduction in total moved tonnage coincides with higher waiting around loading and movement, as well as a higher restart penalty.

In addition to production and waiting outcomes, dynamic dispatch is evaluated on adherence to shift-level bench targets. This acts as the operational proxy for following the planned extraction schedule. The compliance summary shows that dynamic dispatch achieves a mean absolute blend deviation of 0.41 pp in Week 1 and 0.48 pp in Week 2, as seen in Table 5.5. Global deviation is retained as a secondary tonnage-based compliance indicator, with values of 3.37% in Week 1 and 1.75% in Week 2. The corresponding shift-by-bench compliance tables in Appendix I (Tables I.11 and I.12) show that blend and tonnage deviations remain generally small at the shift-bench level, although larger tonnage deviations are concentrated in a limited number of shifts.

The dynamic dispatch setting maintains tight target adherence in both weeks, particularly in Week 2, but this comes with a small cost in system throughput and higher waiting at loaders. This trade-off is consistent with the design goal of dynamic dispatch in this thesis, which is to improve schedule compliance while maintaining production rather than maximizing production alone.

The core scenarios show three distinct patterns across both weeks. First, the autonomous LC case increases total moved tonnage strongly and changes the balance between LC and hauler contributions. In this case, higher LC crusher feed is accompanied by reduced hauler crusher feed and increased hauler stockpiling. Second, the autonomous hauler case increases total moved tonnage modestly and primarily through higher hauler crusher feed. Third, dynamic dispatch reduces total moved tonnage only slightly and is accompanied by higher loader and headway waiting. However, it maintains strong blend adherence to the shift-level bench targets, compared to what would be expected from an unconstrained production-maximizing policy.

Table 5.4: Core scenario results. Week 1 scenarios are listed first, followed by Week 2 scenarios. Weekly totals are aggregated over shifts. Waiting metrics and restart penalty are averaged over runs.

Week / Scenario	Total moved	Hauler → crusher	Hauler → stockpile	LC → crusher	Gain (%)	LQ pos	LQ incl0	Headway	Restart
Week 1									
Baseline	154,769.9	92,068.9	34,733.7	27,967.3	0.00	143.36	50.66	18.08	1.35
DynamicDispatch	153,983.0	91,798.1	32,787.9	28,174.9	-0.51	205.36	57.80	24.04	2.15
Auto_Hauler	157,116.5	93,374.9	35,772.2	27,969.4	+1.52	135.91	44.18	17.02	1.35
Auto_LG54	155,884.6	92,384.6	35,441.3	28,058.7	+0.72	143.78	50.95	17.96	1.41
Auto_LC53	166,580.7	84,401.2	42,417.7	39,761.8	+7.63	144.60	49.31	17.36	1.39
Week 2									
Baseline	151,309.0	101,205.6	26,882.6	23,220.8	0.00	133.46	44.06	23.47	1.41
DynamicDispatch	150,577.8	100,057.9	24,648.0	23,038.4	-0.48	192.75	54.81	28.58	1.74
Auto_Hauler	154,041.5	103,434.9	27,702.4	22,904.2	+1.81	131.45	39.52	18.33	1.43
Auto_LG54	152,270.5	101,907.1	27,317.3	23,046.1	+0.64	136.57	44.39	22.74	1.46
Auto_LC53	163,811.6	95,412.6	32,994.1	35,404.9	+8.26	132.94	42.26	23.67	1.41

LQ pos = average queue at loader (positive waits only, s). LQ incl0 = average queue at loader (including zeros, s). Headway = average headway wait (positive waits only, s). Restart = restart penalty KPI (as reported by the aggregator).

Table 5.5: Dynamic dispatch compliance summary. Mean absolute blend dev is reported as the primary fraction-based compliance metric, with Global dev retained as a tonnage-based compliance indicator.

Week	Mean abs blend dev (pp)	Global dev (%)	Total target (t)	Total abs dev (t)
Week 1 (generalization)	0.41	3.37	129,702	4,370.42
Week 2 (calibration)	0.48	1.75	127,854	2,241.60

5.6.1. Additional scenarios

As outlined in Chapter 4 the additional experiments serve to isolate the impact of specific operational constraints in the core scenario set. The results of the three additional scenarios are presented in the following sections.

Scaling autonomy with three autonomous haulers

Scaling autonomy to three haulers increases total moved in both weeks and reduces average waiting during steady operation. In Week 1, mean total moved increases from 154,769.9 t in the Baseline to 160,210.3 t, a 3.52% increase. In Week 2, mean total moved increases from 151,309.0 t to 156,819.3 t, a 3.64% increase.

The restart penalty also increases in the three hauler case compared to the Baseline configuration. This indicates that, although overall production increases, the restart periods remain associated with more concentrated waiting than in the nominal configuration.

WT sweep for dispatch trade off behavior

The WT sweep quantifies the throughput compliance trade off that is controlled by the target weight parameter. Across both weeks, increasing WT reduces mean absolute blend deviation but also reduces mean tonnage. The effect is strongest at the lower end of the WT values. When WT is near zero, mean absolute blend deviation is highest, indicating that the dispatch policy places limited emphasis on the shift level bench targets. Once WT becomes positive, mean absolute blend deviation drops sharply and then remains within a relatively narrow band for the remainder of the sweep. The figures depicting the trade-off between compliance and throughput are depicted in Appendix I in Figure I.14.

Mean tonnage shows the opposite trend. As WT increases, mean tonnage declines gradually in both weeks. The throughput reduction is largest between WT = 0 and the first few positive WT values, after which the curve flattens.

Global deviation is retained as a secondary compliance indicator in the appendix tables, but the main WT sweep figure now reports mean absolute blend deviation (pp) as the primary compliance metric.

System resilience under one hauler removed

Removing one hauler reduces total moved under both fixed assignment and dynamic dispatch. Under fixed assignment, total tonnage is 141,839.6 t in Week 1 and 137,402.9 t in Week 2. Under dynamic dispatch, total tonnage is 141,641.4 t in Week 1 and 131,058.1 t in Week 2. Compared to fixed assignment under the same failure, dynamic dispatch therefore produces slightly less tonnage in Week 1 and noticeably less tonnage in Week 2.

The compliance results differ by metric. In terms of global deviation, fixed assignment yields 16.51% in Week 1 and 15.57% in Week 2, while dynamic dispatch yields 13.08% and 16.48%, respectively. In terms of mean blend deviation, fixed assignment yields 6.08 pp in Week 1 and 6.70 pp in Week 2, while dynamic dispatch yields 0.43 pp and 0.65 pp. Table 5.6 summarizes these results.

The corresponding blend compliance tables in Appendix I (Tables I.13 and I.15) show that fixed assignment produces much larger shift-bench blend errors than dynamic dispatch. In both weeks, dynamic dispatch generally remains close to the target mix values, while fixed assignment includes several substantially larger deviations at the shift-bench level.

Table 5.6: System resilience results with one hauler removed from the schedule. Dynamic Dispatch and Fixed Dispatch are compared on total tonnage and compliance metrics.

Week	Scenario	Total tonnage (t)	Global dev (%)	Mean blend dev (pp)
Week 1	DynamicDispatch	141,641.4	13.08	0.43
Week 1	FixedDispatch	141,839.6	16.51	6.08
Week 2	DynamicDispatch	131,058.1	16.48	0.65
Week 2	FixedDispatch	137,402.9	15.57	6.70

6

Discussion

6.1. Purpose and reading guide

This chapter interprets the results presented in Chapter 5 in terms of the operational mechanisms of the Flandersbach system. The objective is not to introduce new experimental evidence, but to explain why the tested dispatching and autonomy configurations affect throughput, waiting times, and schedule adherence in the manner observed. The discussion first revisits the variability–flexibility framing introduced earlier in the thesis, then interprets the DD and autonomy results, including the additional scenarios and discusses the influence of the assumed squeezing magnitude implemented in the autonomy experiments. Then, the limitations of this research are discussed, which form the basis of further research.

6.2. Variability and flexibility as interpretation lens

Chapters 2 and 3 place mine performance in the context of variability and flexibility. This variability arises from stochastic execution of tasks, heterogeneous driving behavior, and equipment or personnel unavailability. Flexibility is created through redundancy, routing alternatives, and assignment decisions that can redistribute limited capacity.

DD primarily acts as a flexibility mechanism. It reallocates haulage capacity in response to the current system state and target deviations. It does not remove variability from the system, rather it is a mechanism to counteract the negative impact of this variability.

Autonomy primarily acts as a variability reduction mechanism. By reducing behavioral variance in driving and operational execution and by removing human induced interruptions, autonomy decreases stochastic fluctuations, allowing for a more continuous operation. In specific subsystems, such as LC, autonomy can also increase effective capacity with a single autonomous unit when it reduces non-productive time.

6.3. Dynamic Dispatch

The comparison of the core scenarios in Table 5.4 shows that DD slightly reduces total moved tonnage, while increasing loader queuing, headway waiting, and the restart penalty. At the same time, adherence to the bench targets remains strong in the core DD scenario. In this setting, target adherence is therefore obtained at the cost of a small reduction in production and more waiting around loading and restart periods.

The WT sweep in Section 5.6.1 shows this trade-off. As shown in Figure I.14, increasing WT reduces mean absolute blend deviation while gradually lowering mean total moved tonnage. The steepest reduction in deviation occurs when moving from WT near zero to low positive values, after which the improvement becomes smaller while throughput continues to decline. A useful reference point is WT = 0.5. At this setting, Week 1 throughput is about 156.5 kt for DD compared with about 154.7 kt for the

fixed assignment baseline. In Week 2, DD reaches about 151.5 kt compared to 151.3 kt for the fixed baseline. At the same time, the mean absolute blend deviation is around 1.4 to 1.5 pp. This suggests that relatively low WT values can keep deviations small while increasing the production numbers

This result helps to place the core DD scenario in context. The core setting represents one relatively strict choice on the trade-off curve. The WT sweep shows that DD should not be interpreted as a single fixed policy with a single fixed outcome. Low WT values leave more room for production, while higher values keep the realized mix closer to the target at the cost of throughput. When DD is implemented the WT setting should therefore be kept as low as possible, while maintaining a small enough deviation from the optimal blend to maximize the gains in production.

However, changing WT also changes the optimization setting under which the dispatch behavior was evaluated. Therefore the full consequences of each WT setting are not yet fully understood. A separate retuning step for a fixed WT value could potentially improve production further, but it could also weaken target adherence. This remains a relevant direction for further research.

The scenario with one hauler removed from the schedule highlights where DD is most valuable. Fixed assignment produces significantly larger local deviations at benches where the removed hauler was assigned to previously. Because of the flexibility of DD, this dispatch strategy keeps the realized mix much closer to the targets. The benefit of DD therefore appears most strongly when the system is no longer operating under ideal conditions.

This means that DD does not improve operational stability uniformly. It improves stability in the sense of target adherence, but worsens it in terms of queue formation, headway interactions, and restart concentration in the core scenario.

6.4. Autonomy, coupling, and bottlenecks

The autonomy experiments show that the system-level value of variance reduction depends strongly on where the autonomous unit operates in the production chain.

In the autonomous LC scenario, total moved increases substantially relative to the Baseline, as reported in Table 5.4. The gain is 7.63% in Week 1 and 8.26% in Week 2. This improvement stems solely from increased crusher feed from the LC, this means that in this configuration the increased presence of the LC does not reduce the production of the hauler fleet. The main reason this scenario yields the best results is that LC can operate independently in contrast to the hauler fleet. It is therefore the only vehicle type for which autonomy clearly increases utilization, and this higher utilization translates directly into higher tonnage.

In contrast, introducing a single autonomous hauler yields only modest gains. Total moved increases by 1.52% in Week 1 and 1.81% in Week 2 relative to the Baseline. Although the autonomous hauler operates with reduced behavioral variance, it remains present in the same road network and is loaded by the same loaders as the conventional fleet. Consequently, one autonomous hauler does not create a strong system wide synergy, since the surrounding haulers still introduce variability into the transport flow. The waiting metrics show some improvement, but the production gain remains limited. This suggests that the more continuous transport pattern of one autonomous hauler is still diluted by the surrounding conventional fleet and the shared bottlenecks of the system.

The autonomous loader case also yields only modest gains, with total moved increasing by 0.72% in Week 1 and 0.64% in Week 2. This suggests that reducing variability at the loader improves performance only slightly when the broader system remains constrained by shared interactions between loaders, haulers, and dumping points.

The scaled autonomy scenario with three autonomous haulers shows that increasing the number of autonomous haulers improves throughput more clearly than the single-hauler case. The gain is 3.52% in Week 1 and 3.64% in Week 2. Loader queuing is also lower than in the single-hauler case in both weeks, and headway waiting is lower in Week 1. This is consistent with the idea that when several autonomous haulers assigned to the same part of the system follow the same more continuous behavior, the benefit is influenced less strongly by the conventional fleet. Therefore in this setup the system wide synergy is partially achieved. At the same time, the restart penalty rises strongly. During breaks and

shift changes these haulers wait at the locations from which they will be loaded. When operations resume, three haulers are already queued at the loading point, which increases the restart penalty.

These results help explain why the autonomy scenarios differ so much across subsystems. LC benefits most because autonomy increases its effective utilization in a part of the system that can operate relatively independently. A single autonomous hauler benefits less because it remains embedded in the same shared network as the conventional fleet. A larger number of autonomous haulers does create more synergy, because the haulers are assigned to the same bench and can benefit from the reduced variance of the other haulers without being hindered by the conventional fleet in large portions of their cycles.

Additionally, the effect of the squeeze parameter was evaluated. These checks do not change the ranking of the core scenarios, and the influence of the parameter remains limited. Across the tested values of 2, 3, 5, and 6, no clear pattern emerges as can be seen in Tables J.1 and J.2. Although the squeeze parameter pulls sampled values toward the mean, the more extreme settings do not produce systematically better or worse performance. This suggests that the model is not sensitive to the exact assumed squeeze magnitude. A likely explanation is that the parameter only affects one vehicle, and any local reduction in variability is diluted by the remaining system dynamics.

TThe autonomy experiments therefore show that reducing variability at a single point in the system does not automatically improve operational stability at system level. The effect depends on how independent the subsystem is and on how strongly the remaining manual fleet continues to shape the overall outcomes.

6.5. Context dependence and external validity

In Figure 2.1 differences in AHS adoption between regions is depicted. These differences likely stem from the operational context of these regions. Flandersbach is a small mine compared to the large scale operations in Australia and Canada, limiting the traveled distance in the haulage cycle and increasing the amount of interactions. These conditions make Flandersbach, and European mines in general, less favorable for AHS. Larger operations often run continuously, which gives more hours over which the fixed cost of AHS can be recovered. Longer haul roads also contain longer and more repetitive driving segments. These are generally easier for AHS than dense interaction zones around loaders, crushers, and intersections.

Economic context is likely to influence the adoption of AHS too. In Australia and Canada mines are both remote and operate under harsher conditions compared to Germany. This changes the economics of automation. In these remote operations, companies make use of airplanes to transport personnel and provide on-site housing in extreme conditions. These costs make AHS attractive in more cases than in a smaller European operation. Additionally the salaries of personnel in Australia are higher compared to Germany, reducing the incentive for autonomous options even further Salary Expert, 2024; Stuff Limited, 2024.

The DES used in this thesis is configured specifically for Flandersbach, however the model structure supports wider deployment. The simulation architecture is not tied to a fixed layout or a hard-coded set of parameters, but is built around modular submodels as can be seen from Figure E.1 and Figure H.1. Most operational inputs are sampled from external datafiles, following specific preprocessing steps. Other inputs were obtained via interviews with the mine personnel. Changing these input files and other mine specific parameters results in the representation of a different operation, provided that the required data files are available and are in line with the code configuration. This does not mean that the model can be transferred without caution. The current implementation is calibrated to the operational context of Flandersbach In another mine conditions can differ substantially, which can change both the behavior of the system and the relevance of specific submodels or parameters. The framework should therefore be recalibrated before being used for quantitative assessment elsewhere. Furthermore, the model has not been tested on a different mining setup yet, meaning that unforeseen errors and failure points may still emerge when the model is used in a different setup. The same reasoning applies to commodity type. The model could be used for a different commodity than limestone, but adjustments may be required when ore grades, blend constraints, or destination-specific processing logic play a more dominant role.

The LC result should also be generalized carefully. Its strength in this case does not mean that LC autonomy will always outperform haul truck autonomy. The broader lesson is that autonomy creates the most value where the autonomous unit can make use of capacity that would otherwise be underused.

The broader conclusion is therefore that the percentage effects reported in this thesis should not be transferred directly to other mines, but that the DES framework and the comparative logic of the interventions are transferable when the model is recalibrated to local infrastructure, fleet behavior, schedules, and operational constraints.

6.6. Limitations affecting interpretation

Several modeling and data limitations influence the interpretation of the results.

First, the transportation model relies on a simplified speed formulation rather than empirical segment level speed recordings. While the speed validation tables in Appendix I.5 and Appendix I.6 demonstrate reasonable consistency, implementing empirical speed values per road segment would reduce reliance on a heuristic equation and likely improve the quality of the DES.

Second, weather effects were analyzed using averaged data from two nearby weather stations. The plotted diagnostics include mean daily temperature, mean daily snow, mean daily precipitation, cumulative precipitation over the previous seven days, mean daily frost, hours during operation with temperature below zero, mean frost night per day, and indicators of sub-zero night temperatures. Plots of these variables are shown in Figure J.3 in Appendix J. Although the influence of these variables were explored, the current model does not couple weather conditions to the transportation model component. These variables are known to influence the achievable speeds significantly in mines, because of their influence on the road quality. Incorporating these variables would improve the speed model of the DES and create the opportunity for better production estimates.

Third, maintenance data was analyzed in terms of failure durations and time between failures, as shown in Figure J.1 in Appendix J. These metrics are directly relevant for stochastic breakdown modeling. In the present implementation maintenance and equipment failure is not incorporated. A more detailed integration of empirical failure distributions would allow the model to capture this variability, and test the performance of the scenarios under these circumstances. This would likely result in different results for the DD experiment, as this experiment has shown and is designed to counteract such disturbances.

Fourth, empirical inputs were trimmed to remove sensor artifacts and extreme outliers. While necessary for robustness, this reduces the frequency of rare but operationally relevant disruptions. The resulting bias likely leads to optimistic performance estimates and may understate the value of the DD flexibility mechanisms that are most beneficial under disruptive conditions. Secondly these removed disruptions are important for the autonomy experiments too. With an autonomous unit these disruptions could be removed, therefore including them in the baseline setup could result in more significant production increases in the autonomy scenarios.

Fifth, human variability is represented indirectly. Personal performance data are not recorded, however operational teams remain constant. On even weeks, one team operates early and late shifts, and on odd weeks the shift allocation reverses. Differences in operational efficiency between weeks are visible at group level without. Examples of these differences are illustrated in Figure J.2 in Appendix J. Incorporating such structured group-level variability would improve realism compared to a uniform variability factor.

These limitations and the results suggest that the model is most useful for comparing scenarios under a common set of assumptions, instead of providing exact production estimates. However hauler production appears to be overestimated systematically showing comparable deviations in both weeks. This bias could in principle be used to scale the simulated hauler outputs toward more realistic levels. The overestimation of hauler production can be explained by model simplifications. In the DES haulers never wait at the crushers for other haulers, whereas they have to do so in reality. Loader relocations during shifts has not been modeled either. Loaders travel at low speeds, making the time these relocations take significant. During such relocations all assigned haulers are unable to haul material. Leaving these relocations and the waiting for other haulers out of the DES therefore partially explains the observed increase in production. The LC results are less systematic in this respect, likely because

the share of time spent on auxiliary tasks was not known. If that task split becomes available, it could be incorporated into the model to improve the LC representation as well.

7

Conclusion

The objective of this research was to quantify the operational impact of DD and three distinct autonomous vehicle options at the Flandersbach limestone mine. To do so, this thesis developed and calibrated a DES digital twin using empirical inputs derived from telemetry data and implementations of the operational rules of the mine. Then a set of scenarios designed to answer the research questions were evaluated over two representative operating weeks. This chapter summarizes the main findings to answer the research questions defined in Chapter 1.

The first research question assessed whether the heuristic DD approach implemented in this thesis is an effective method to increase production volume. However across both simulated weeks, the DD scenario reduced total moved tonnage relative to the fixed-assignment Baseline, while increasing congestion indicators, most notably loader queue time and headway waiting (Table 5.4). Within the scope of this thesis, DD was designed and tuned to account for shift-level bench targets, and the reported compliance metrics show that the policy enforces tight adherence under that definition (Table 5.5 and Appendix I, Tables I.11 and I.12). The observed trade-off is therefore that the current DD implementation functions as a target adherence policy, but does so at a throughput cost driven by increased waiting times and herding behavior. At the same time the WT sweep experiments show that this trade-off is tunable. When the target adherence constraint is loosened the DD strategy can increase production while keeping the target deviations limited. The scenario with a removed hauler shows that the DD strategy that the value of this strategy becomes clearer in a disturbed system, which is in line with the proposed benefits of a more flexible setup.

The second research question evaluated the introduction of single autonomous agents into a manual fleet. The simulation results indicate that the value of a single autonomous hauler or a single autonomous loader is limited in a mixed-fleet setting because productive time remains bounded by the operating schedule of the manual system (Table 5.4). In contrast, the autonomous LC scenario produced the largest throughput gains of all core scenario set in both weeks. This increase coincides with the fact that the LC intervention expands the effective operating window and therefore adds productive capacity. At the same time, this result should be interpreted in light of the modeling scope: the LC in the DES is represented as a dedicated crusher-feeding unit, whereas in reality it also performs auxiliary duties, which is reflected in the validation discussion of subsystem bias (Section 5.4.4).

The third research question considered the operational trade-offs across interventions. In the studied setting, the interventions do not provide uniform improvements in production performance and operational stability across the system. DD changes where and when trucks accumulate, which can increase waiting times and restart effects and reduce throughput. A single autonomous vehicle reduces variability for one asset, but most benefits are absorbed by downstream constraints, because the rest of the system remains manual. In contrast, the LC autonomy case increases productive time, producing the strongest throughput gains. However the autonomous LC represents the intervention with the most demanding deployment requirements as the LC is located in the most interaction-dense area of the operation. When the share of autonomous vehicles is increased, some system wide synergy can be

achieved, as shown in the three autonomous hauler scenario. With a larger quantity of autonomous haulers additional production time can be realized too, by offsetting the loader breaks from each other. The benefits of autonomous vehicles thus scale with the extent of adoption of these vehicles, and the constraint of one autonomous vehicle reduces the impact of the autonomous vehicles in the experimental setting.

Beyond the scenario results themselves, the development and validation of the DES is an additional contribution. The validation results show a consistent positive bias in total moved and total crushed relative to site records, but this bias is stable in direction across both modeled weeks and is traceable to specific modeled subsystems and scope choices (Section 5.4.4 and Appendix I). As a result, the model is suitable as a decision-support tool for scenario comparisons under consistent assumptions.

In conclusion, the findings suggest that replacing individual manual units with autonomous ones yields limited returns when the remainder of the haulage cycle remains constrained by shared resources, the default schedules, and interaction effects. The largest modeled gains arise when an intervention increases effective operating time. For future exploration, the results point toward configurations that change the constraints currently limiting continuous operation. This could be realized by either increasing the autonomy share or by redesigning certain operating procedures. For DD the core experiments show a slight underperformance compared to the fixed assignment baseline. The potential of DD is more visible in the additional experiments. The single hauler failure experiment represents a realistic operational disturbance and shows that in this scenario DD can maintain the mine plan more effectively. The WT sweep experiments indicate that the DD algorithm can increase production compared to the baseline for low WT values, but at the cost of larger deviations from the targets. This indicates that the value of this DD approach strongly depends on the selection of an acceptable range of these deviations.

References

Abolghasemian, M., Daneshmand-Mehr, M., & Ghane, A. (2020). A two-phase simulation-based optimization of hauling system in open-pit mine. *Iranian Journal of Management Studies*, *13*, 705–732. https://doi.org/10.22059/ijms.2020.294809.673898

Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. https://arxiv.org/abs/1907.10902

Alarie, S., & Gamache, M. (2002). Overview of solution strategies used in truck dispatching systems for open pit mines. *International Journal of Surface Mining, Reclamation and Environment*, *16*(1), 59–76.

Aoshima, K. (2025). High-performance autonomous wheel loading: A computational approach. https://api.semanticscholar.org/CorpusID:275569059

Backman, S., Lindmark, D., Bodin, K., Servin, M., Mörk, J., & Löfgren, H. (2021). Continuous control of an underground loader using deep reinforcement learning. *Machines*, *9*(10), 216. https://doi.org/10.3390/machines9100216

Bakhtavar, E., & Mahmoudi, H. (2020). Development of a scenario-based robust model for the optimal truck-shovel allocation in open-pit mining. *Computers & Operations Research*, *115*. https://doi.org/10.1016/j.cor.2018.08.003

Barde, S., Yacout, S., & Shin, H. (2019). Optimal preventive maintenance policy based on reinforcement learning of a fleet of military trucks. *Journal of Intelligent Manufacturing*, *30*. https://doi.org/10.1007/s10845-016-1237-7

Bergstra, J., Bardenet, R., Bengio, Y., & Kégl, B. (2011). Algorithms for hyper-parameter optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, & K. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 24). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf

Bernardi, L., Kumral, M., & Renaud, M. (2020). Comparison of fixed and mobile in-pit crushing and conveying and truck-shovel systems used in mineral industries through discrete-event simulation. *Simulation Modelling Practice and Theory*, *103*, 102100. https://doi.org/10.1016/j.simpat.2020.102100

Burgess-Limerick, R., Horberry, T., Lynas, D., Hill, A., & Haight, J. (2024, December). Human aspects of mining automation. https://doi.org/10.1201/9781003380887-1

Caterpillar Inc. (2000, April). *Cat 777D Off-Highway Truck: Technical Specifications* [Document ID: AEHQ5121-02. Applicable to vehicles SKW 16, 17]. Caterpillar Inc. https://www.cat.com/

Caterpillar Inc. (2013, February). *Cat 777G Off-Highway Truck: Technical Specifications* [Document ID: AEHQ6822-01. Applicable to vehicles SKW 28, 29, 32, 36, 37]. Caterpillar Inc. https://www.cat.com/

Caterpillar Inc. (2022, May). *Cat 777 (07) Off-Highway Truck Bare Chassis: Technical Specifications* [Document ID: CM20220523-f3e1b-b8388. Applicable to vehicle SKW 40]. Caterpillar Inc. https://www.cat.com/

Caterpillar Inc. (n.d.). *Cat 992 Wheel Loader: Product Specifications* [Document in Dutch (Productspecificaties). Applicable to vehicle LG 55]. Caterpillar Inc. https://www.cat.com/

Chaowasakoo, P., Seppälä, H., Koivo, H., & Zhou, Q. (2017). Improving fleet management in mines: The benefit of heterogeneous match factor. *European Journal of Operational Research*. https://doi.org/10.1016/j.ejor.2017.02.039

Corke, P., Roberts, J., Cunningham, J., & Hainsworth, D. (2008, January). Mining robotics. https://doi.org/10.1007/978-3-540-30301-5_50

Dadhich, S. (2018). *Optimization of surface mining operations using dynamic dispatch and data analytics* [Doctoral Thesis]. Luleå University of Technology [Printed by Luleå University of Technology, Graphic Production 2018].

Dadhich, S., Bodin, U., & Andersson, U. (2016). Key challenges in automation of earth-moving machines. *Automation in Construction*, *68*, 212–222. https://doi.org/10.1016/j.autcon.2016.05.009

Dadhich, S., Sandin, F., Bodin, U., Andersson, U., & Martinsson, T. (2019). Field test of neural-network based automatic bucket-filling algorithm for wheel-loaders. *Automation in Construction*, *97*, 1–12. https://doi.org/https://doi.org/10.1016/j.autcon.2018.10.013

David Bird. (2019, May). *Autonomous mining equipment* (Report 2 - New Technology & Innovation No. 2). RFC Ambrian. https://www.google.com/url?sa=t&source=web&rct=j&opi=89978449&url=https://www.rfcambrian.com/wp-content/uploads/2019/04/RFCA-NTI-Report-2-Autonomous-Mining-Equipment-May-2019.pdf&ved=2ahUKEwj3vs7M84WRAxUBOfsDHWEdGXoQFnoECBwQAQ&usg=AOvVaw0A512IvZilcVeKpPiwtUI2

de Carvalho, J. P., & Dimitrakopoulos, R. (2021). Integrating production planning with truck-dispatching decisions through reinforcement learning while managing uncertainty. *Minerals*, *11*(6). https://doi.org/10.3390/min11060587

Du, H., Chan, L., Tong, J., Raad, R., Naghdy, F., Guo, Q., Yu, Y., Islam, M. R., Tubbal, F., Ros, M., Li, Z., & Ritz, C. (2025). Industrial progress of robotic automation in mining applications: A survey. *Mining, Metallurgy & Exploration*, *42*, 537–556. https://doi.org/10.1007/s42461-025-01219-y

Equipment Journal. (2021, October). Volvo announces lx03 electric autonomous wheel loader concept. https://www.equipmentjournal.com/construction-equipment/volvo-announces-lx03-electric-autonomous-wheel-loader-concept/

Eriksson, D., Ghabcheloo, R., & Geimer, M. (2024). Optimizing bucket-filling strategies for wheel loaders inside a dream environment. *Automation in Construction*, *168*, 105804. https://doi.org/https://doi.org/10.1016/j.autcon.2024.105804

Fernando, H., & Marshall, J. (2020). What lies beneath: Material classification for autonomous excavators using proprioceptive force sensing and machine learning. *Automation in Construction*, *119*, 103374. https://doi.org/10.1016/j.autcon.2020.103374

Firoozi, A. A., Tshambane, M., Firoozi, A. A., & Sheikh, S. M. (2024). Strategic load management: Enhancing eco-efficiency in mining operations through automated technologies. *Results in Engineering*, *24*, 102890. https://doi.org/https://doi.org/10.1016/j.rineng.2024.102890

Gaber, T., El Jazouli, Y., Eldesouky, E., & Ali, A. (2021). Autonomous haulage systems in the mining industry: Cybersecurity, communication and safety issues and challenges. *Electronics*, *10*(11). https://doi.org/10.3390/electronics10111357

Global Mining Guidelines Group. (2024, July). *Guideline for the implementation of autonomous systems in mining* (tech. rep.) (Version 2). Global Mining Guidelines Group. https://gmggroup.org/wp-content/uploads/2024/08/GUIDELINE_Implementation-of-Autonomous-Systems.pdf

Goli, M., Moniri-Morad, A., Aguilar, M., Shishvan, M. S., Shahsavar, M., & Sattarvand, J. (2025). A simulation-based risk assessment model for comparative analysis of collisions in autonomous and non-autonomous haulage trucks. *Applied Sciences*, *15*(17), 9702. https://doi.org/10.3390/app15179702

Gustafson, A. (2011). Dependability assurance for automatic load haul dump machines. https://api.semanticscholar.org/CorpusID:115276103

Hitachi Construction Machinery Co., Ltd. (2018, June). *Hitachi EX1200-7 Hydraulic Excavator: Technical Specifications* [Document ID: KS-EN408. Applicable to vehicle LG 14]. Hitachi Construction Machinery Co., Ltd. https://www.hitachi-c-m.com/

Icarte, G., Berrios, P., Castillo, R., & Herzog, O. (2020). A multiagent system for truck dispatching in open-pit mines. In M. Freitag, H.-D. Haasis, H. Kotzab, & J. Pannek (Eds.), *Dynamics in logistics: Proceedings of the 7th international conference ldic 2020, bremen, germany* (pp. 363–373). Springer Nature Switzerland AG. https://doi.org/10.1007/978-3-030-44783-0

Icarte, G., Pinto, J., & Herzog, O. (2021, March). A dynamic scheduling multiagent system for truck dispatching in open-pit mines. https://doi.org/10.1007/978-3-030-71158-0_6

Icarte-Ahumada, G., & Herzog, O. (2025). Intelligent scheduling in open-pit mining: A multi-agent system with reinforcement learning. *Machines*, *13*(5). https://doi.org/10.3390/machines13050350

Iqbal, A., Joshi, Agrawal, S., & Bakhsh, F. (2021, January). *Metaheuristic and evolutionary computation: Algorithms and applications*. https://doi.org/10.1007/978-981-15-7571-6

Kaur, D. (2024). The impact of autonomous vehicles on mining operations: Enhancing safety and productivity through technological advancements. *Scholarly Review Journal*, *SR Online: Showcase*. https://doi.org/10.70121/001c.124875

Komatsu Europe International N.V. (2018, December). *Komatsu HD785-7 Off-Highway Truck: Technical Specifications* [Document ID: EENSS20341. Applicable to vehicles SKW 38, 39]. Komatsu Europe International N.V. https://www.komatsu.eu/

Liebherr-Werk Bischofshofen GmbH. (2025, February). Innovations in liebherr wheel loaders: Autonomous operation and hydrogen drive [Accessed: 22 November 2025].

Long, M., Schafrik, S., Kolapo, P., Agioutantis, Z., & Sottile, J. (2024). Equipment and operations automation in mining: A review. *Machines*, *12*(10). https://doi.org/10.3390/machines12100713

Md-Nor, Z. A., Kecojevic, V., Komljenovic, D., & Groves, W. (2008). Risk assessment for haul truck-related fatalities in mining [Society for Mining, Metallurgy & Exploration (SME)]. *Mining Engineering*, *60*(3), 44–49.

Meech, J., & Parreira, J. (2011). An interactive simulation model of human drivers to study autonomous haulage trucks. *Procedia CS*, *6*, 118–123. https://doi.org/10.1016/j.procs.2011.08.023

Mena, R., Zio, E., Kristjanpoller, F., & Arata, A. (2013). Availability-based simulation and optimization modeling framework for open-pit mine truck allocation under dynamic constraints. *International Journal of Mining Science and Technology*, *23*(1), 113–119. https://doi.org/https://doi.org/10.1016/j.ijmst.2013.01.017

Meng, S., Tian, B., Zhang, X., Qi, S., Zhang, C., & Zhang, Q. (2024). Openmines: A light and comprehensive mining simulation environment for truck dispatching, 3043–3049. https://doi.org/10.1109/IV55156.2024.10588764

MiningTechnology. (2022). Australia continues to dominate the use of autonomous haul trucks.

Mirzaei-Nasirabad, H., Mohtasham, M., Askari Nasab, H., & Alizadeh, B. (2023). An optimization model for the real-time truck dispatching problem in open-pit mining operations. *Optimization and Engineering*, *24*. https://doi.org/10.1007/s11081-022-09780-x

Mohtasham, M., Mirzaei-Nasirabad, H., Askari-Nasab, H., & Alizadeh, B. (2022). Multi-stage optimization framework for the real-time truck decision problem in open-pit mines: A case study on sungun copper mine. *International Journal of Mining, Reclamation and Environment*, *36*(7), 461–491. https://doi.org/10.1080/17480930.2022.2067709

Moniri-Morad, A., Shishvan, M. S., Aguilar, M., Goli, M., & Sattarvand, J. (2024). Powered haulage safety, challenges, analysis, and solutions in the mining industry; a comprehensive review. *Results in Engineering*, *21*, 101684. https://doi.org/https://doi.org/10.1016/j.rineng.2023.101684

Moradi, A., & Askari Nasab, H. (2017). Mining fleet management systems: A review of models and algorithms. *International Journal of Mining, Reclamation and Environment*, *33*, 1–19. https://doi.org/10.1080/17480930.2017.1336607

Moradi Afrapoli, A., Tabesh, M., & Askari-Nasab, H. (2017). *An investigation into dispatch optimizers using truck-shovel simulation and a new multi objective truck dispatching technique* (tech. rep. No. MOL Report Eight 201). Mining Optimization Laboratory (MOL), University of Alberta.

Munirathinam, M., & Yingling, J. C. (1994). A review of computer-based truck dispatching strategies for surface mining operations. *International Journal of Surface Mining, Reclamation and Environment*, *8*(1), 1–15.

Newman, A., Rubio, E., Caro, R., Weintraub, A., & Eurek, K. (2010). A review of operations research in mine planning. *Informs Journal of Applied Analytics*, *40*, 222–245. https://doi.org/10.1287/inte.1090.0492

Noriega, R. (2023). *An optimization and machine learning framework for dynamic dispatching in mining operations* [PhD thesis]. University of Alberta. https://doi.org/10.7939/r3-03sh-0w11

Noriega, R., & Pourrahimian, Y. (2022). A systematic review of artificial intelligence and data-driven approaches in strategic open-pit mine planning. *Resources Policy*, *77*, 102727. https://doi.org/https://doi.org/10.1016/j.resourpol.2022.102727

Noriega, R., Pourrahimian, Y., & Askari-Nasab, H. (2025). Deep reinforcement learning based real-time open-pit mining truck dispatching system. *Computers & Operations Research*, *173*, 106815. https://doi.org/https://doi.org/10.1016/j.cor.2024.106815

Ozdemir, B., & Kumral, M. (2019). Simulation-based optimization of truck-shovel material handling systems in multi-pit surface mines. *Simulation Modelling Practice and Theory*, *95*, 36–48. https://doi.org/https://doi.org/10.1016/j.simpat.2019.04.006

Parreira, J. (2013). *I an interactive simulation model to compare an autonomous haulage truck system with a manually-operated system* [Doctoral dissertation, The Univeristy of British Columbia].

Peres, F., & Castelli, M. (2021). Combinatorial optimization problems and metaheuristics: Review, challenges, design, and development. *Applied Sciences*, *11*(14). https://doi.org/10.3390/app11146449

Rogers, W., Kahraman, M., Drews, F., Powell, K., Haight, J., Wang, Y., Baxla, K., & Sobalkar, M. (2019). Automation in the mining industry: Review of technology, systems, human factors, and political risk. *Mining, Metallurgy & Exploration*, *36*. https://doi.org/10.1007/s42461-019-0094-2

Salary Expert. (2024). Truck driver dump truck salary germany [Accessed: 2024-10-24]. https://www.salaryexpert.com/salary/job/truck-driver-dump-truck/germany

Saleem, H. A. (2025). Automation and artificial intelligence in enhancing mining efficiency and sustainability: A review. *Procedia Environmental Science, Engineering and Management*, *12*(1), 213–228. http://www.procedia-esem.eu

Smith, A., Linderoth, J., & Luedtke, J. (2021). Optimization-based dispatching policies for open-pit mining. *Optimization and Engineering*, *22*, 1–41. https://doi.org/10.1007/s11081-021-09628-w

Soleymani Shishvan, M., & Benndorf, J. (2018). Simulation-based optimization approach for material dispatching in continuous mining systems. *European Journal of Operational Research*, *275*. https://doi.org/10.1016/j.ejor.2018.12.015

Soofastaei, A., Aminossadati, S., Kizil, M., & Knights, P. (2016). A discrete-event model to simulate the effect of truck bunching due to payload variance on cycle time, hauled mine materials and fuel consumption. *International Journal of Mining Science and Technology*, *26*(5), 745–752. https://doi.org/https://doi.org/10.1016/j.ijmst.2016.05.047

Stewart, L. (2022, December). Caterpillar partners with a customer to bring robot haul trucks to smaller sites [Image provided by Caterpillar. Accessed: March 1, 2024]. https://www.constructionbriefing.com/news/caterpillar-partners-with-a-customer-to-bring-robot-haul-trucks-to-smaller-sites/8025520.article

Stuff Limited. (2024, October). Truck driver in australian mine reveals staggering salary for just 6 months work [Accessed: 2024-10-24]. https://www.stuff.co.nz/money/350475533/truck-driver-in-australian-mine-reveals-staggering-salary-for-just-6-months-work

Subtil, R., Silva, D., & Alves, J. (2011). A practical approach to truck dispatch for open pit mines. *Proceedings of the 35th APCOM Symposium*, 1–10.

Tampier, C., Mascaró, M., & Ruiz-del-Solar, J. (2021). Autonomous loading system for load-haul-dump (lhd) machines used in underground mining. *Applied Sciences*, *11*(18). https://doi.org/10.3390/app11188718

Temeng, V. A., Otuonye, F. O., & Frendewey Jr, J. O. (1997). Real-time truck dispatching using a transportation algorithm. *International Journal of Surface Mining, Reclamation and Environment*, *11*(4), 203–207.

Voronov, A. Y., Dubinkin, D. M., & Voronov, Y. (2023). An overview of models for truck dispatching in open-pit mines. *Mining Industry Journal (Gornay Promishlennost)*. https://api.semanticscholar.org/CorpusID:255255337

Voronov, Y., Voronov, A., & Makhambayev, D. (2020). Current state and development prospects of autonomous haulage at surface mines. *E3S Web of Conferences*, *174*, 01028. https://doi.org/10.1051/e3sconf/202017401028

Wang, X., Dai, Q., Bian, Y., Xie, G., Xu, B., & Yang, Z. (2023). Real-time truck dispatching in open-pit mines. *International Journal of Mining, Reclamation and Environment*, *37*, 1–20. https://doi.org/10.1080/17480930.2023.2201120

White, J. W., & Olson, J. P. (1986). Computer-based dispatching in mines with concurrent operating objectives. *Mining Engineering*, *38*(11), 1045–1054.

Zeppelin Cat. (n.d.). *Cat 992K Wheel Loader: Technical Specifications* [Accessed: 2026-01-02. Specifies bucket volume 10.7–12.3 m^3. Applicable to vehicles LG 53, 54]. Zeppelin Nederland B.V. https://www.zeppelin-cat.nl/nl/producten/wielladers/grote-wielladers/cat-992k

Zhang, L., Song, X., Wang, L., Qian, L., Guan, T., He, Z., Ye, Z., Song, R., Ding, H., & Manocha, D. (2021). Autonomous excavator system with real-world deployment.

A

Hauler Data

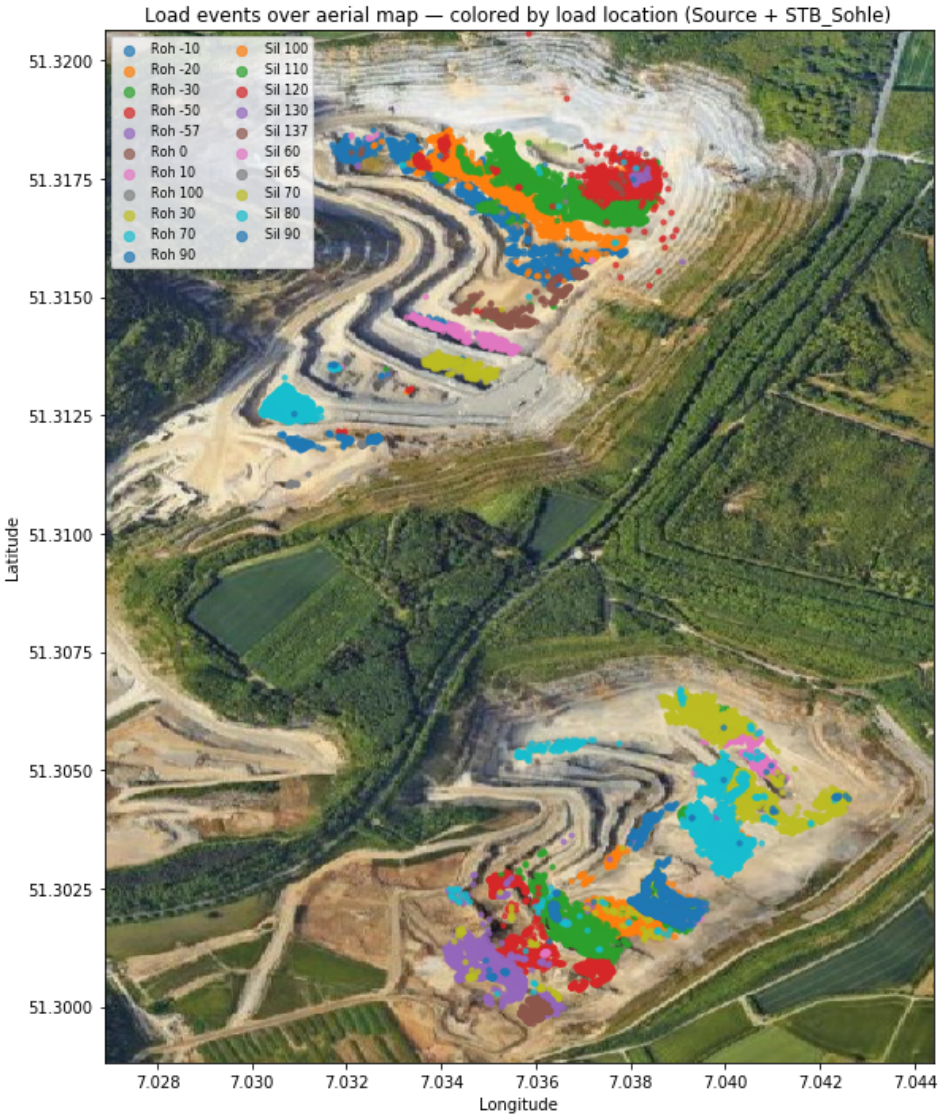

Figure A.1: Aerial map with dots corresponding to different benches assigned via criteria.

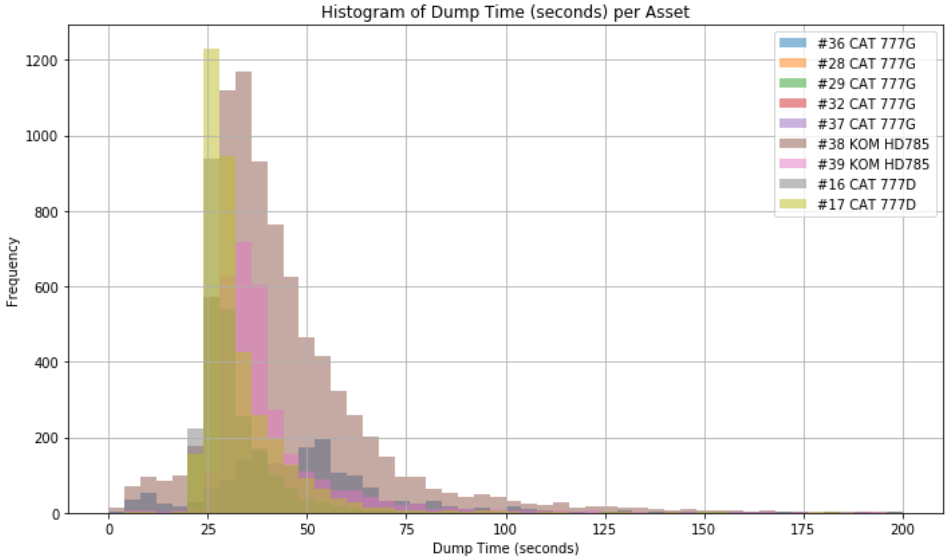

Figure A.2: Histogram of hauler dump times.

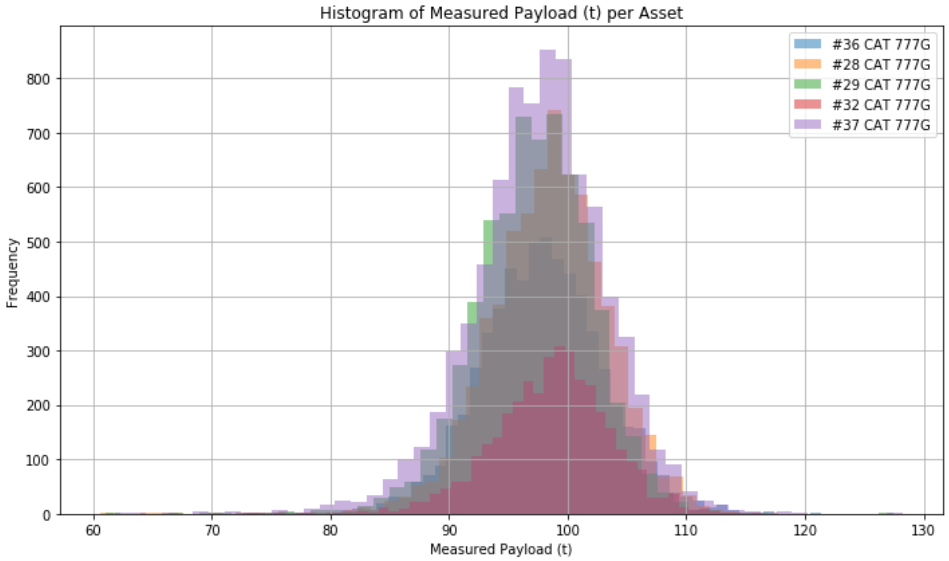

Figure A.3: Histogram of hauler payload (t).

B

Loader

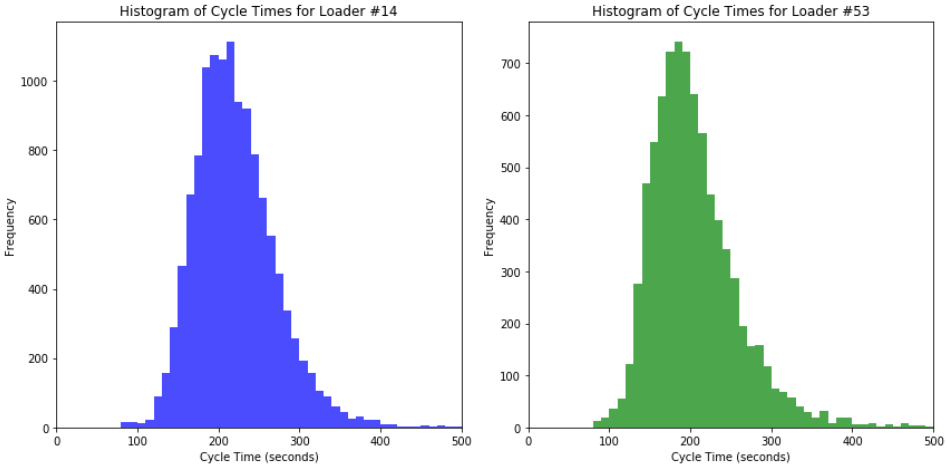

Figure B.1: Histograms of full cycle times of loader 14 and loader 53.

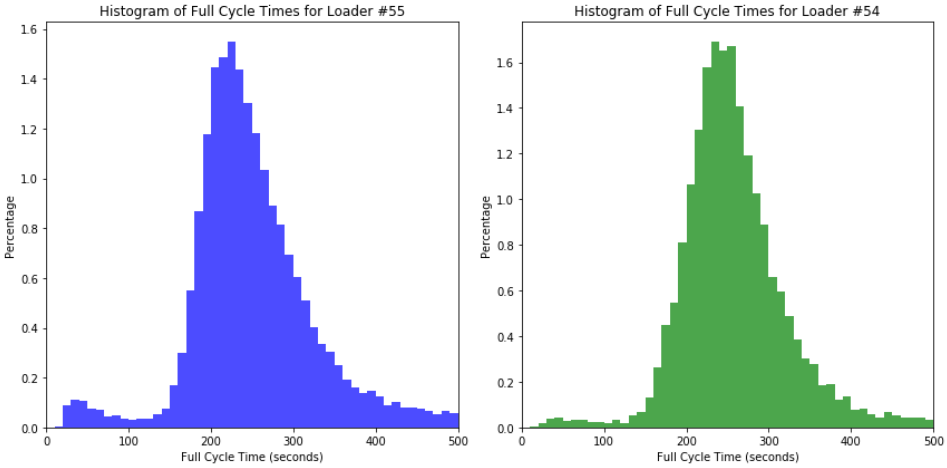

Figure B.2: Histograms of full cycle times of loader 54 and loader 55.

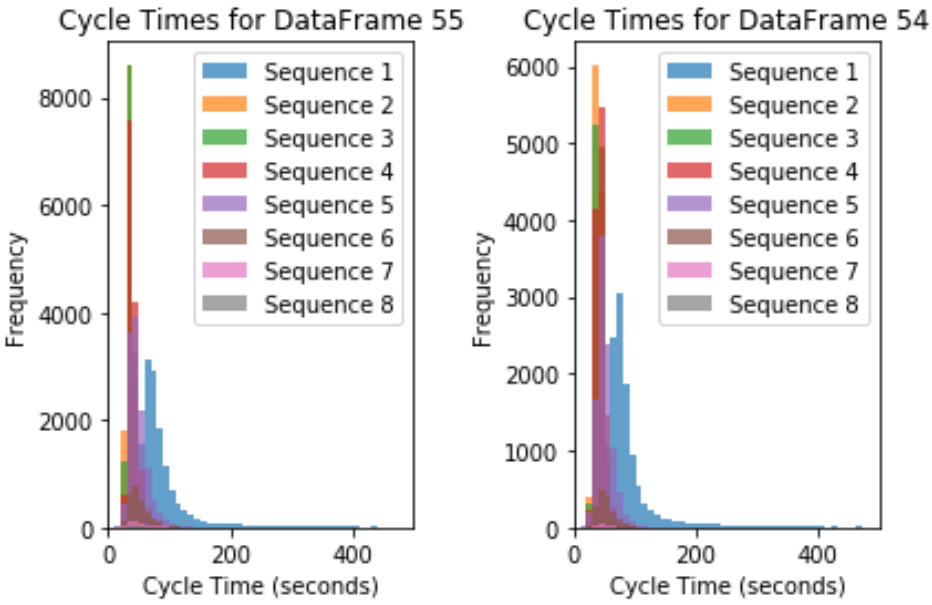

Figure B.3: Histogram of cycle time per bucket payload sequence for loader 54 and loader 55.

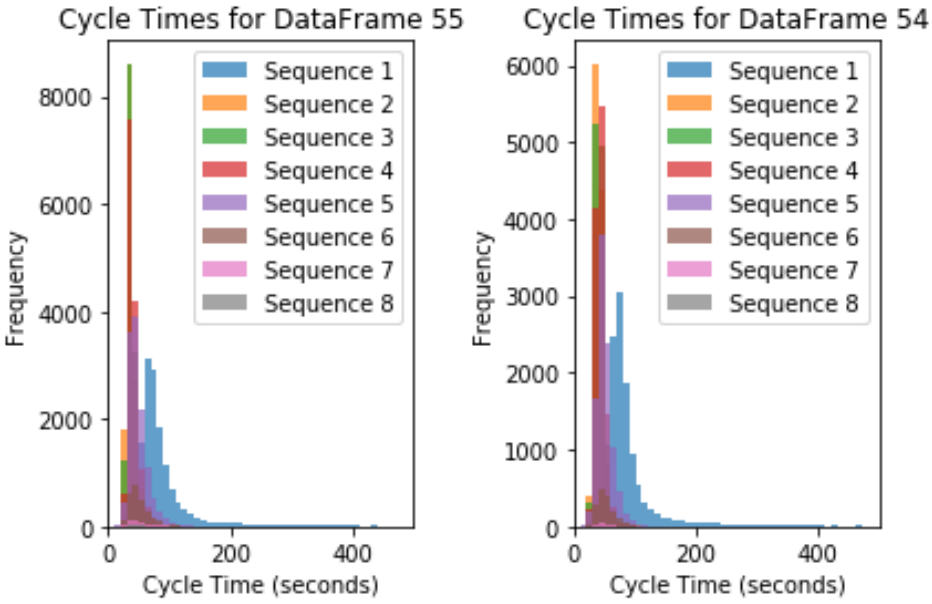

Figure B.4: Histogram of cycle time per bucket payload sequence for loader 54 and loader 55.

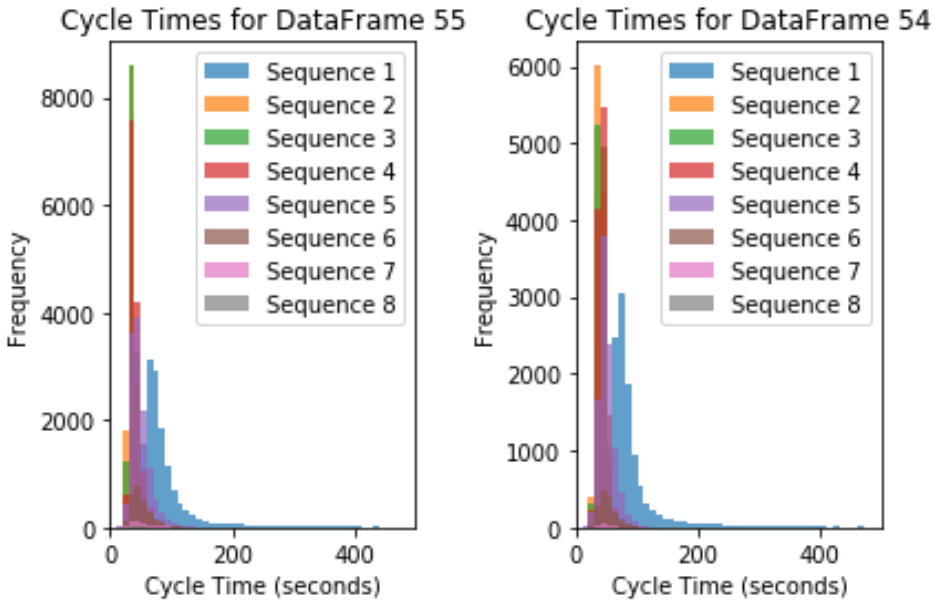

Figure B.5: Histogram of cycle time per bucket payload sequence for loader 54 and loader 55 with exclusion of sequences 1, and 6 or larger.

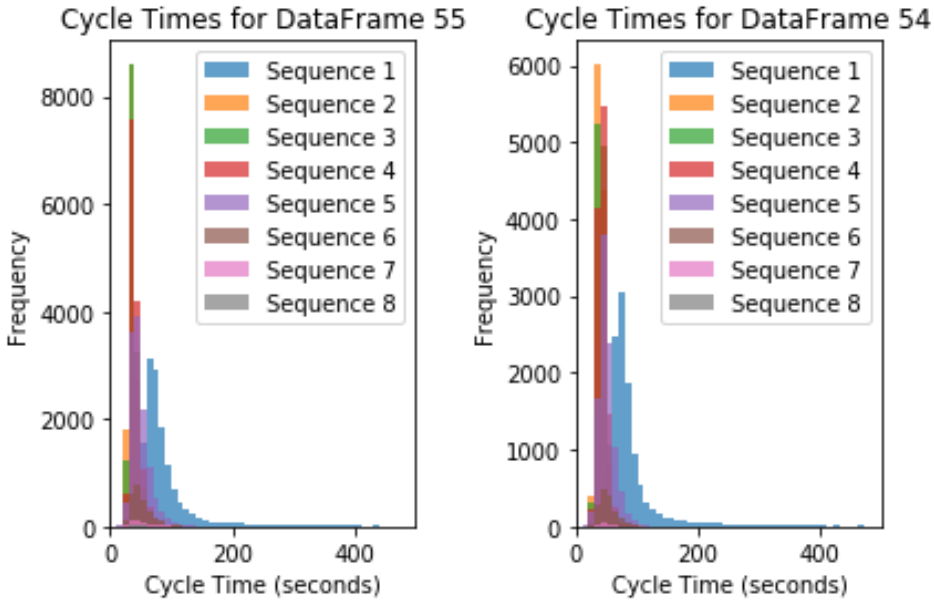

Figure B.6: Histogram of payload per bucket for sequence<6 for loader 54 and loader 55.

C

Load and Carry

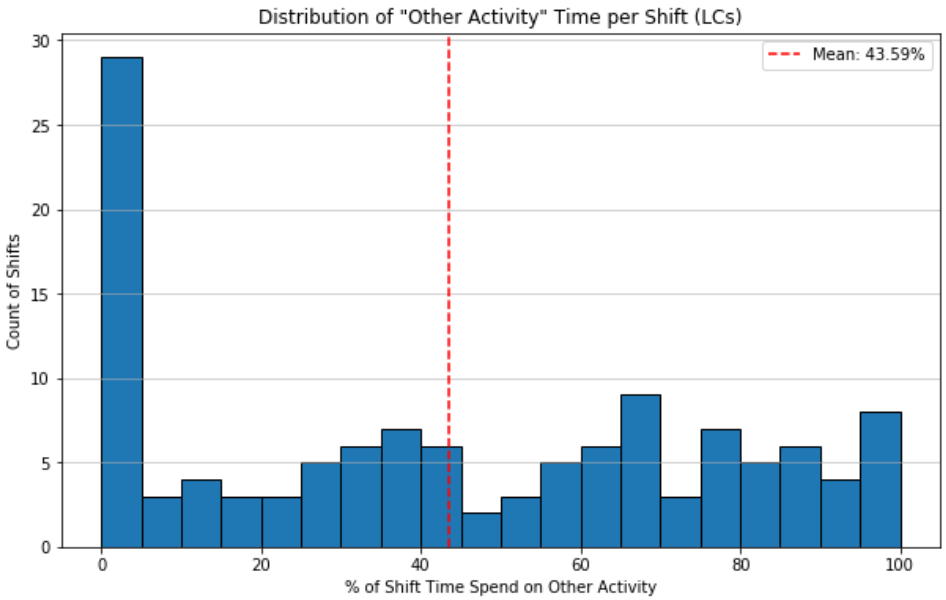

Figure C.1: Histograms showing amount of shifts, with percentage of ”other activity” during the shift.

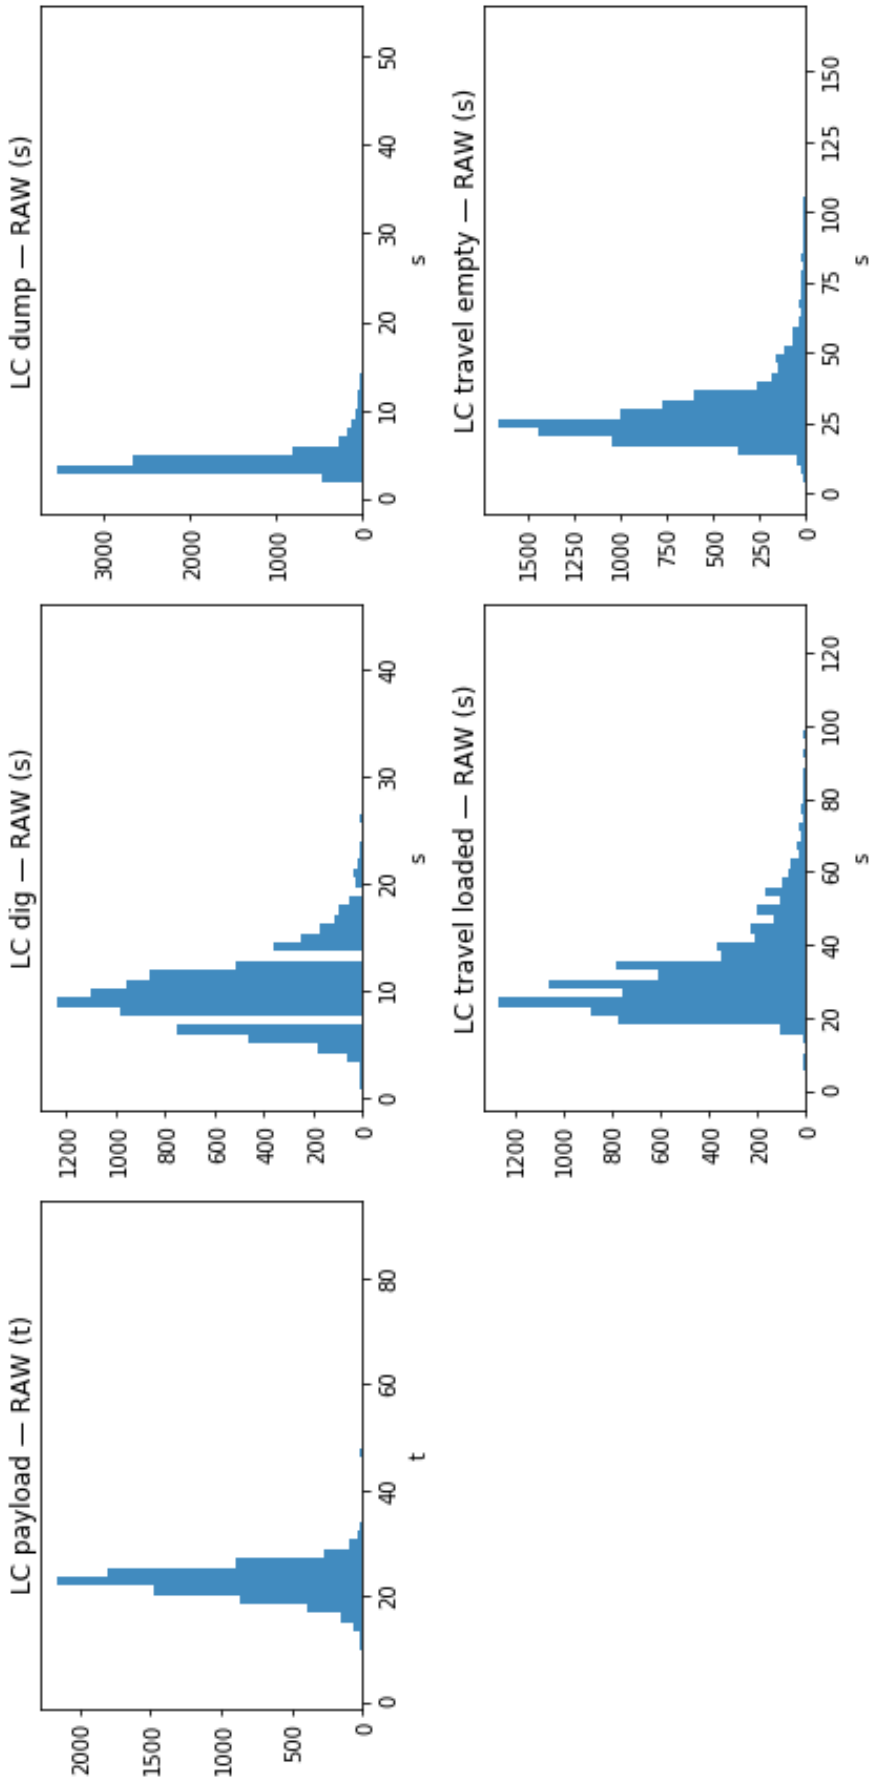

Figure C.2: Distributions of the main activities of the load-and-carry vehicle, aggregated from LC54 and LC55 telemetry data.

D

DES

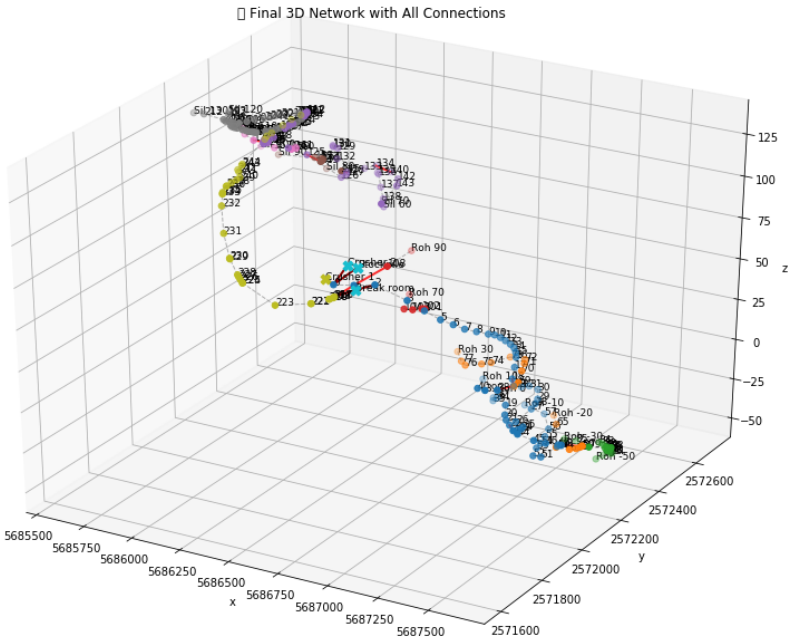

Figure D.1: 3D representation of the road network.

E

Input Validation

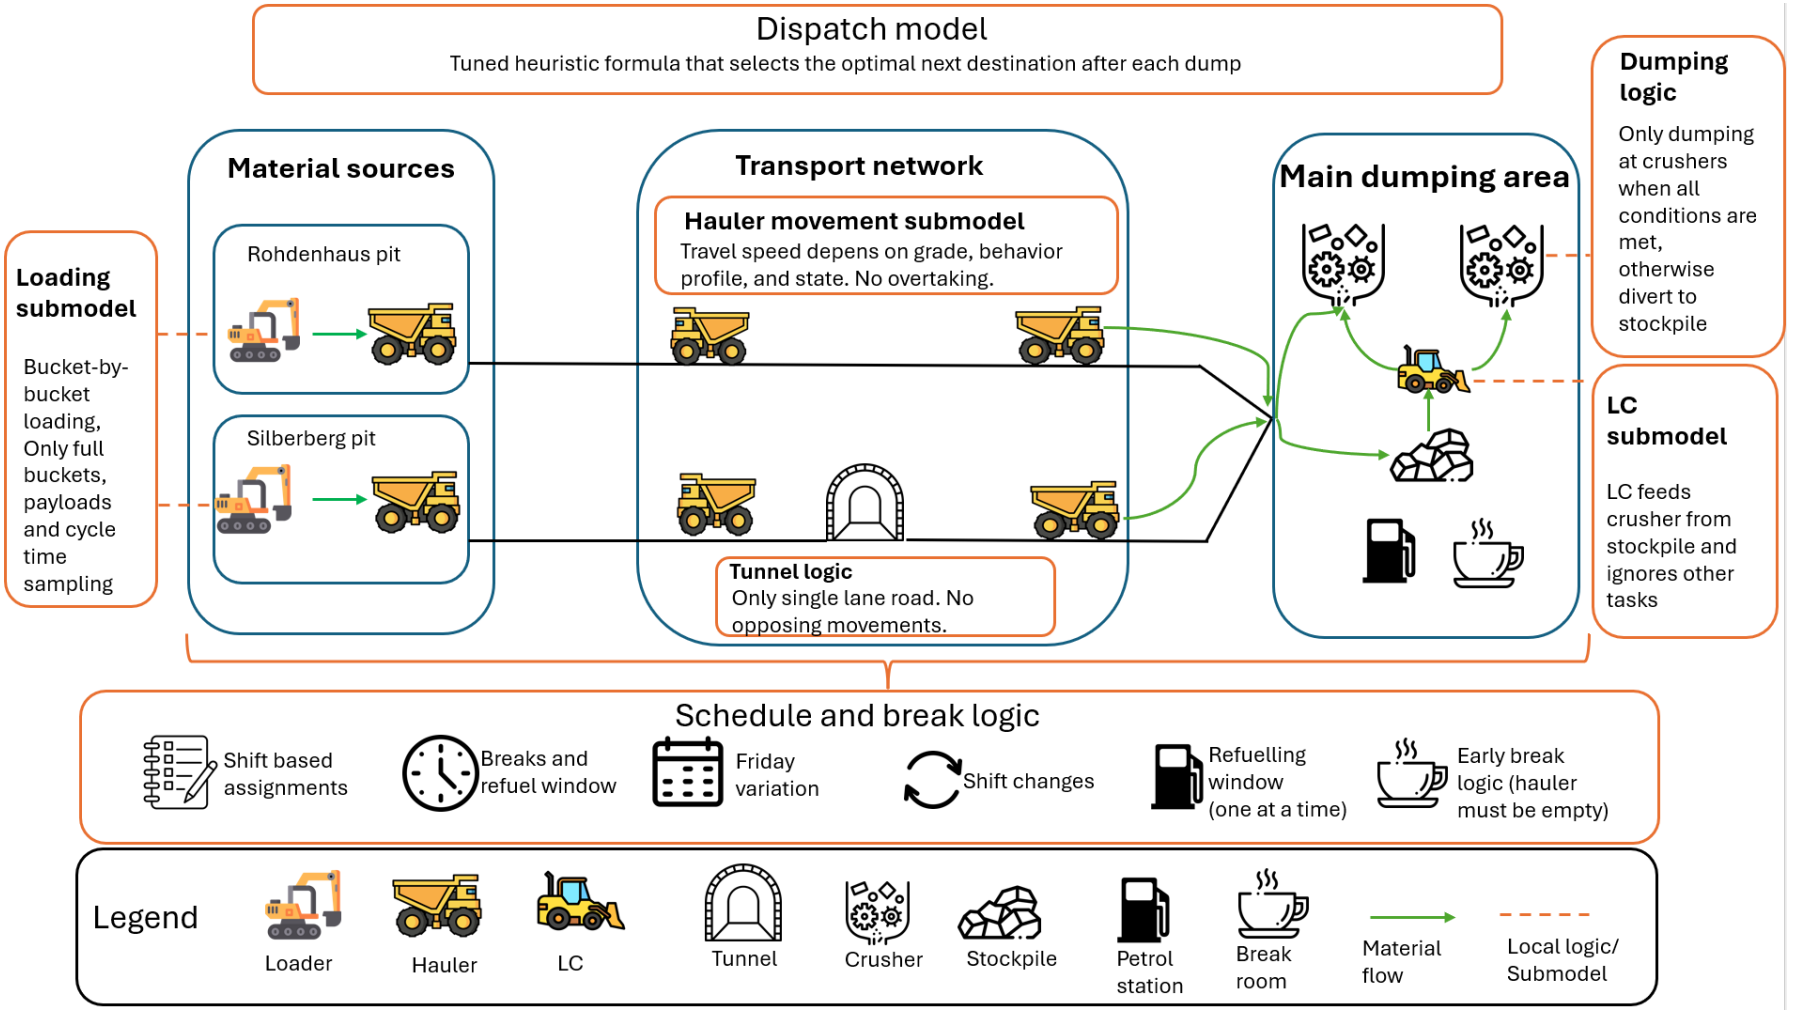

Figure E.1: Schematic overview of the DES model architecture used in this thesis. The figure indicates where the main operational submodels act in the Flandersbach system and distinguishes them from supervisory layers such as dispatch and schedule logic.

Figure E.2: Buffer Level During week 2

F

Company Documents

The following excerpts are taken from internal company documents and are included to document how baseline assignments, shift-level tonnage targets, and reporting totals are derived. The documents use German naming conventions. The shift label F denotes Früh (early shift) and S denotes Spät (late shift). Fleet identifiers follow the site convention, where SKW denotes haul trucks and LG denotes loaders.

Figure F.1 is used to reconstruct baseline dispatch and bench assignment at the shift level. The Loader and Bench columns provide the mapping required to place each loader at a bench. The Hauler and Loader columns specify which haul truck was assigned to which loader within a shift. The Ton column reports the tonnage attributed to each loader–bench–hauler record. These tonnages are aggregated by summing Ton per bench and shift, producing the shift-level totals used for comparisons in Chapter 5. Within-shift loader relocations are rarely observed in the available schedule records. When a loader appears at more than one bench within a shift, the baseline assigns the loader to the bench that accounts for the largest share of recorded Ton during that shift.

Figure F.2 is used to derive shift-level production context for the simulated weeks. The report distinguishes tonnage delivered to the crushers from tonnage routed to the stockpile, and it provides load-and-carry (L&C) reporting at the shift level. In the L&C part of the table, the column labeled Produktion [h] denotes the hours spent operating as load-and-carry during the shift, while Produktion L&C (t) denotes the tonnage fed to the crusher by the load-and-carry unit. The achieved rate is reported as Leistung. Additional totals reported include Brecheraufgabe (crushed material) and Gesamt Bew. Mas. (total moved). In a small number of cases, the reported load-and-carry operating time implies unrealistically high instantaneous rates due to rounding of the operating-hour entry. Where required, the recorded operating time is adjusted by rounding to the nearest full operational hour to keep implied rates within plausible bounds. The variability in implied load-and-carry rates is consistent with the dual role of the unit, which performs crusher feeding as well as auxiliary tasks that reduce effective feeding time.

Date		Shift	Pit	Bench			Loader	Hauler	Destination	Ton	Material Type	
28-2-2025	28-2-2025	F	Rohs.	-50	518		LG 54	SKW 40	BF 22 Br. 185	594	ChO	6
28-2-2025	28-2-2025	F	Rohs.	-50	518		LG 54	SKW 39	BF 22 Br. 185	396	ChO	4
28-2-2025	28-2-2025	F	Rohs.	-50	518		LG 54	SKW 38	BF 22 Br. 185	495	ChO	5
28-2-2025	28-2-2025	F	Rohs.	-20	445		LG 54	SKW 39	BF 22 Br. 185	693	ChO	7
28-2-2025	28-2-2025	F	Rohs.	-20	445		LG 54	SKW 38	BF 22 Br. 185	693	ChO	7
28-2-2025	28-2-2025	F	Rohs.	-20	445		LG 54	SKW 40	BF 22 Br. 185	594	ChO	6
28-2-2025	28-2-2025	F	Sil.	60	2928		LG 14	SKW 29	BF 22 Br. 185	891	ChO	9
28-2-2025	28-2-2025	F	Sil.	60	2928		LG 14	SKW 28	BF 22 Br. 185	792	ChO	8
28-2-2025	28-2-2025	F	Sil.	60	2928		LG 14	SKW 28	BF 22 FL A	99	ChO	1
28-2-2025	28-2-2025	F	Rohs.	90	266	FL A	LG 53	L&C (1)	BF 22 Br. 185	2776	ChO	2776
28-2-2025	28-2-2025	F	Sil.	120	2217		LG 55	SKW 37	BF 22 Br. 185	1188	ChO D	12
28-2-2025	28-2-2025	F	Sil.	120	2217		LG 55	SKW 36	BF 22 Br. 185	1287	ChO	13
28-2-2025	28-2-2025	F	Sil.	120	2217		LG 55	SKW 37	BF 22 FL A	99	ChO	1
28-2-2025	28-2-2025	S	Rohs.	-20	445		LG 54	SKW 38	BF 22 Br. 185	1584	ChO	16
28-2-2025	28-2-2025	S	Rohs.	-20	445		LG 54	SKW 37	BF 22 Br. 185	1485	ChO	15
28-2-2025	28-2-2025	S	Rohs.	-20	445		LG 54	SKW 36	BF 22 Br. 185	1485	ChO	15
28-2-2025	28-2-2025	S	Rohs.	-20	445		LG 54	SKW 36	BF 22 FL A	99	ChO	1
28-2-2025	28-2-2025	S	Sil.	60	2928		LG 14	SKW 28	BF 22 Br. 185	297	ChO	3
28-2-2025	28-2-2025	S	Sil.	60	2928		LG 14	SKW 29	BF 22 Br. 185	297	ChO	3
28-2-2025	28-2-2025	S	Sil.	60	2928		LG 14	SKW 17	BF 22 Br. 185	273	ChO	3
28-2-2025	28-2-2025	S	Sil.	70	2827		LG 14	SKW 17	BF 22 Br. 185	910	ChO	10
28-2-2025	28-2-2025	S	Sil.	70	2827		LG 14	SKW 28	BF 22 Br. 185	1089	ChO	11
28-2-2025	28-2-2025	S	Sil.	70	2827		LG 14	SKW 29	BF 22 Br. 185	990	ChO	10
28-2-2025	28-2-2025	S	Sil.	120	2217		LG 55	SKW 40	BF 22 Br. 185	1485	ChO	15
28-2-2025	28-2-2025	S	Sil.	120	2217		LG 55	SKW 39	BF 22 Br. 185	1485	ChO	15
28-2-2025	28-2-2025	S	Sil.	120	2217		LG 55	SKW 39	BF 22 FL A	99	ChO	1

Figure F.1: Excerpt from the internal full schedule used to reconstruct baseline assignments. Relevant columns include Date, Shift, Pit, Bench, Loader (LG), Hauler (SKW), Destination, and Ton.

		LG, SKW, LC		Produkion SKW			Produkion L&C						
Datum	S	Einsatz-grad Plan	Einsatz-grad Ist	Prod. [h]	Produk-tion [t]	Leistung [t/h]	Prod. [h]	Produk-tion [t]	Leistung [t/h]	Ges. bew.Mas.Eg. [t]	Brecher-aufgabe [t]	SKW FL A zum Br. [t]	AV [%]
24.02.2025	F	3,7,0	3, 7, 1	7.0	11 669		1067.0	1	1067.00	12 736	10 786		
				14.0	**24 823**		**2340.0**	**2**		**27 163**	**22 760**		
25.02.2025	F	3,7,0	4, 7, 1	7.0	12 263		1514.0	2	1009.33	13 777	9 887		
25.02.2025	S	3,7,0	3, 7, 1	7.0	14 080		2750.0	3	916.67	16 830	12 999		
				14.0	**26 343**		**4264.0**	**5**		**30 607**	**22 886**		
26.02.2025	F	3,7,1	4, 7, 1	7.0	10 869		5438.0	6	906.33	16 307	10 283		
26.02.2025	S	3,7,1	4, 7, 1	7.0	11 760		5473.0	7	781.86	17 233	11 760		
				14.0	**22 629**		**10911.0**	**13**		**33 540**	**22 043**		
27.02.2025	F	3,7,0	3, 8, 1	7.0	12 354		228.0	1	456.00	12 582	11 958		
27.02.2025	S	3,7,0	3, 7, 1	7.0	15 613				2230.43	15 613	14 152		
				14.0	**27 967**		**228.0**	**1**		**28 195**	**26 110**		
28.02.2025	S	3,8,0	3, 8, 1	5.8	11 578				2013.57	11 578	11 380		
28.02.2025										1 435			
28.02.2025	F	3,7,1	4, 7, 1	5.8	7 821		2776.0	5	555.20	10 597	7 623		
				11.5	**19 399**		**2776.0**	**5**		**23 610**	**19 003**		

Figure F.2: Excerpt from the internal production report for the simulated weeks. The table reports shift-level production totals, including crusher feed and stockpile-related totals, as well as load-and-carry operating time and crusher-feeding tonnage. German column headers include Brecheraufgabe (crushed material), Gesamt Bew. Mas. (total moved), and Leistung (achieved rate).

G

Visualization

G.1. Visualization example

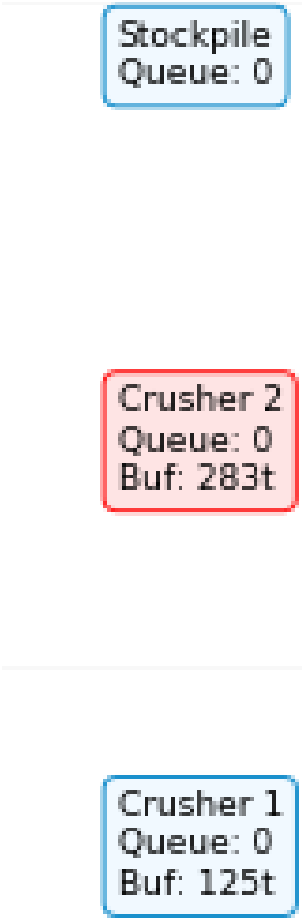

Figure G.1: Main Dumping Area Status of Dump Locations

```
== BENCH STATUS ==
Roh -10     □□ Q:0 W:1
Roh -20        Q:0 W:0
Roh -30        Q:0 W:0
Roh -50        Q:0 W:0
Roh 0          Q:0 W:0
Roh 10         Q:0 W:0
Roh 30         Q:0 W:0
Roh 70         Q:0 W:0
Roh 90      □□ Q:0 W:1
Sil 100        Q:0 W:0
Sil 110        Q:0 W:0
Sil 120        Q:0 W:0
Sil 130        Q:0 W:0
Sil 60         Q:0 W:0
Sil 70         Q:0 W:0
Sil 80         Q:0 W:0
Sil 90         Q:0 W:0
Stockpile      Q:0 W:0
```

Figure G.2: Queuing and Working Status per Bench

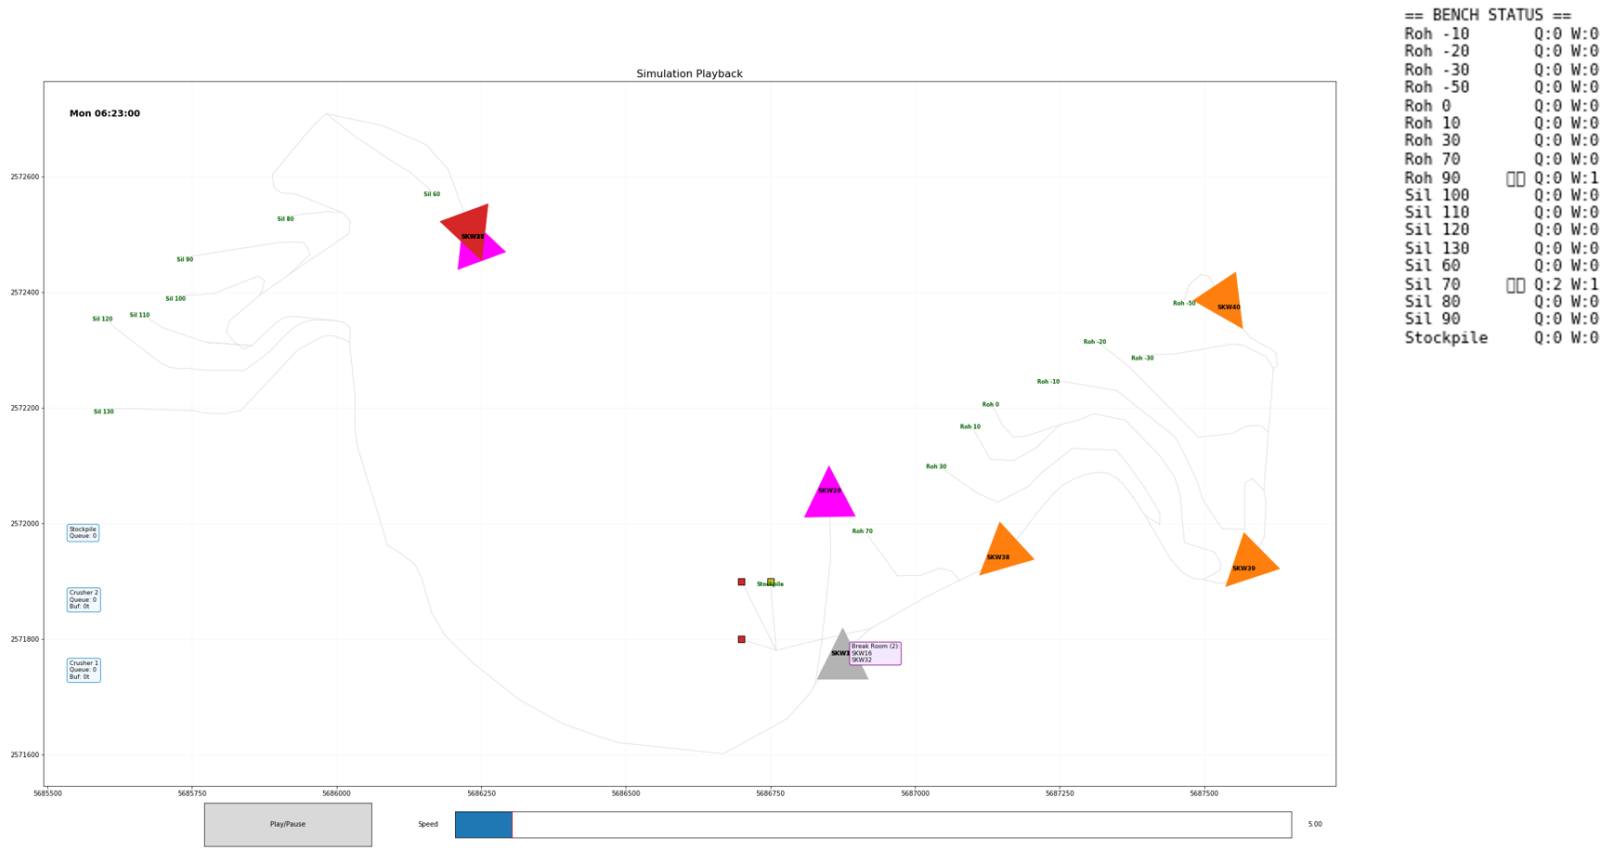

Figure G.3: Playback snapshot at the start of the operational day. The zoomed-in bench-status panel shows that multiple haulers are waiting shortly after the start of the shift. Because the vehicles start at the same time, they arrive at the active benches at approximately the same moment, which creates early waiting and confirms the restart concentration represented in the DES.

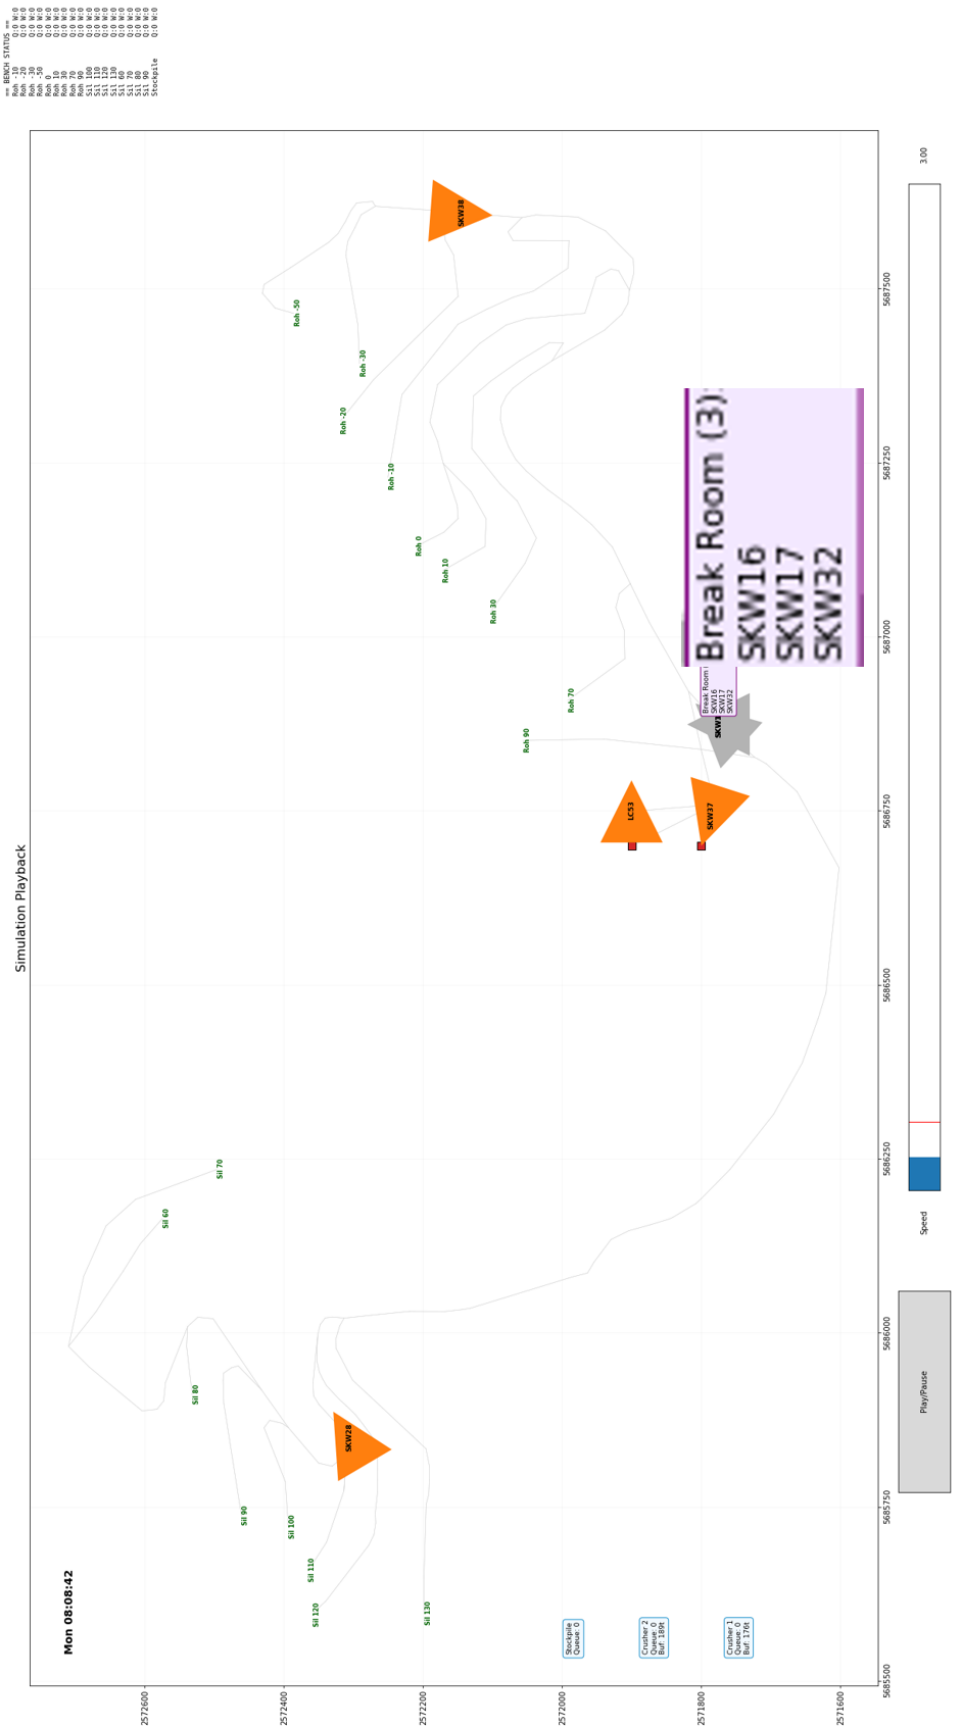

Figure G.4: Playback snapshot during the refuel window at 08:08. Three haulers are visible in the breakroom area: two are parked and one is refueling. This confirms that the fueling rule is executed within the designated refuel window and that the breakroom area also functions as the location for tanking events.

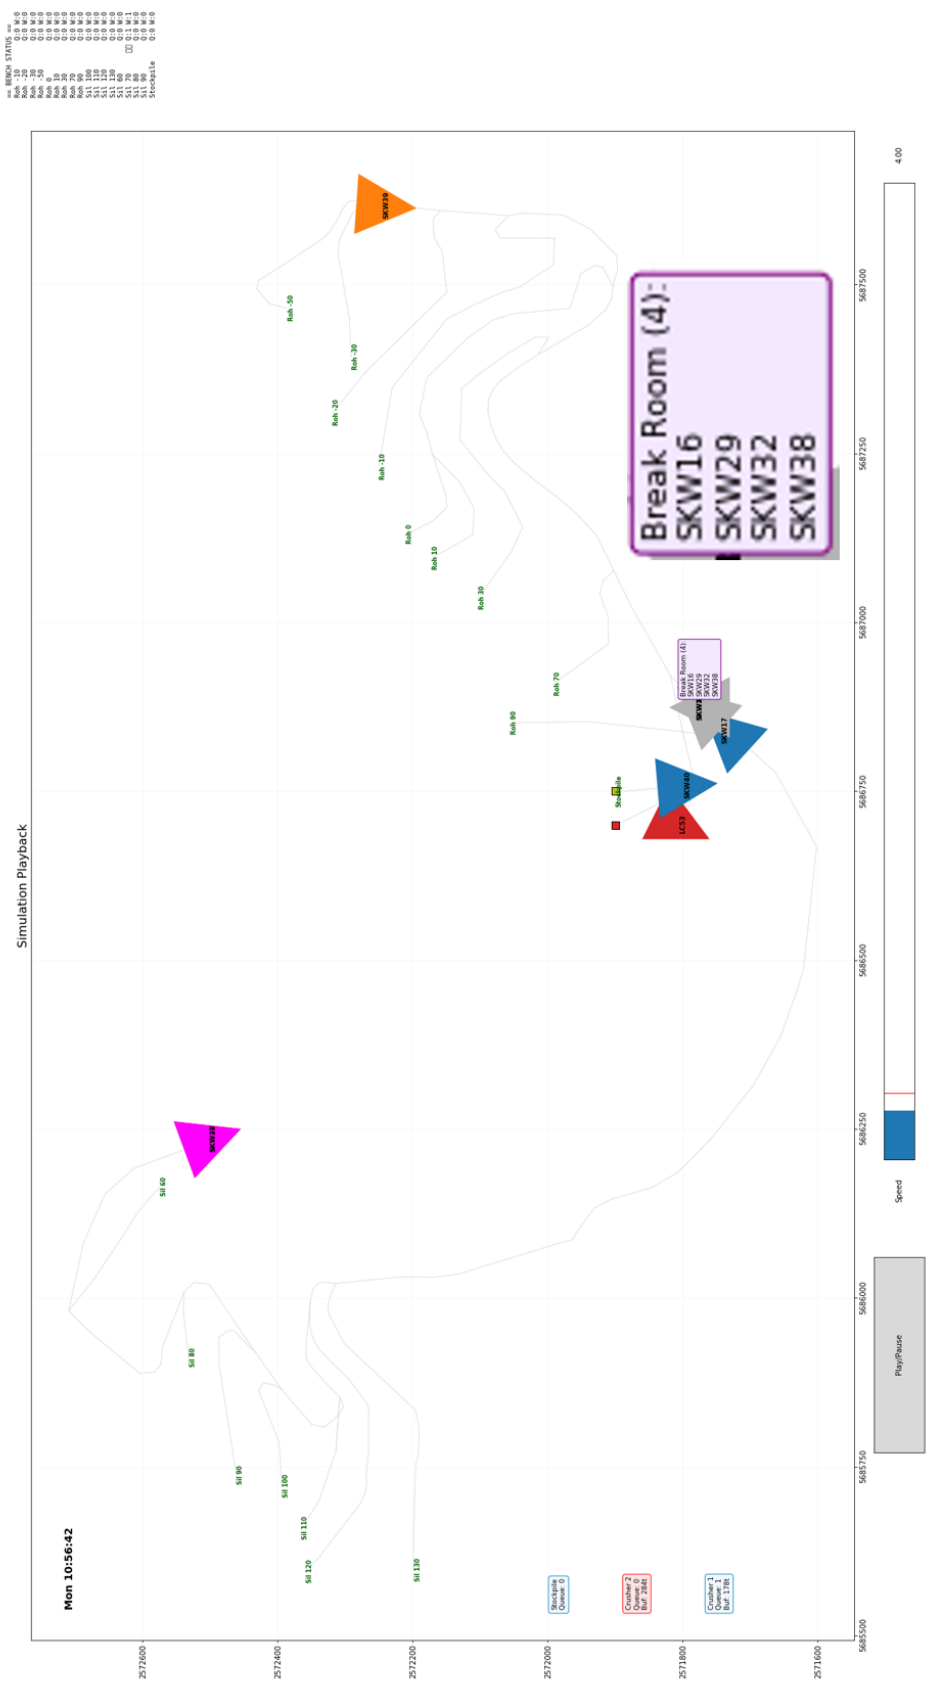

Figure G.5: Playback snapshot four minutes before the scheduled break. The zoomed-in breakroom state shows four haulers in the breakroom: two were already present and two others have entered early. This confirms that the early-break logic is active and allows selected vehicles to enter the breakroom before the official break start.

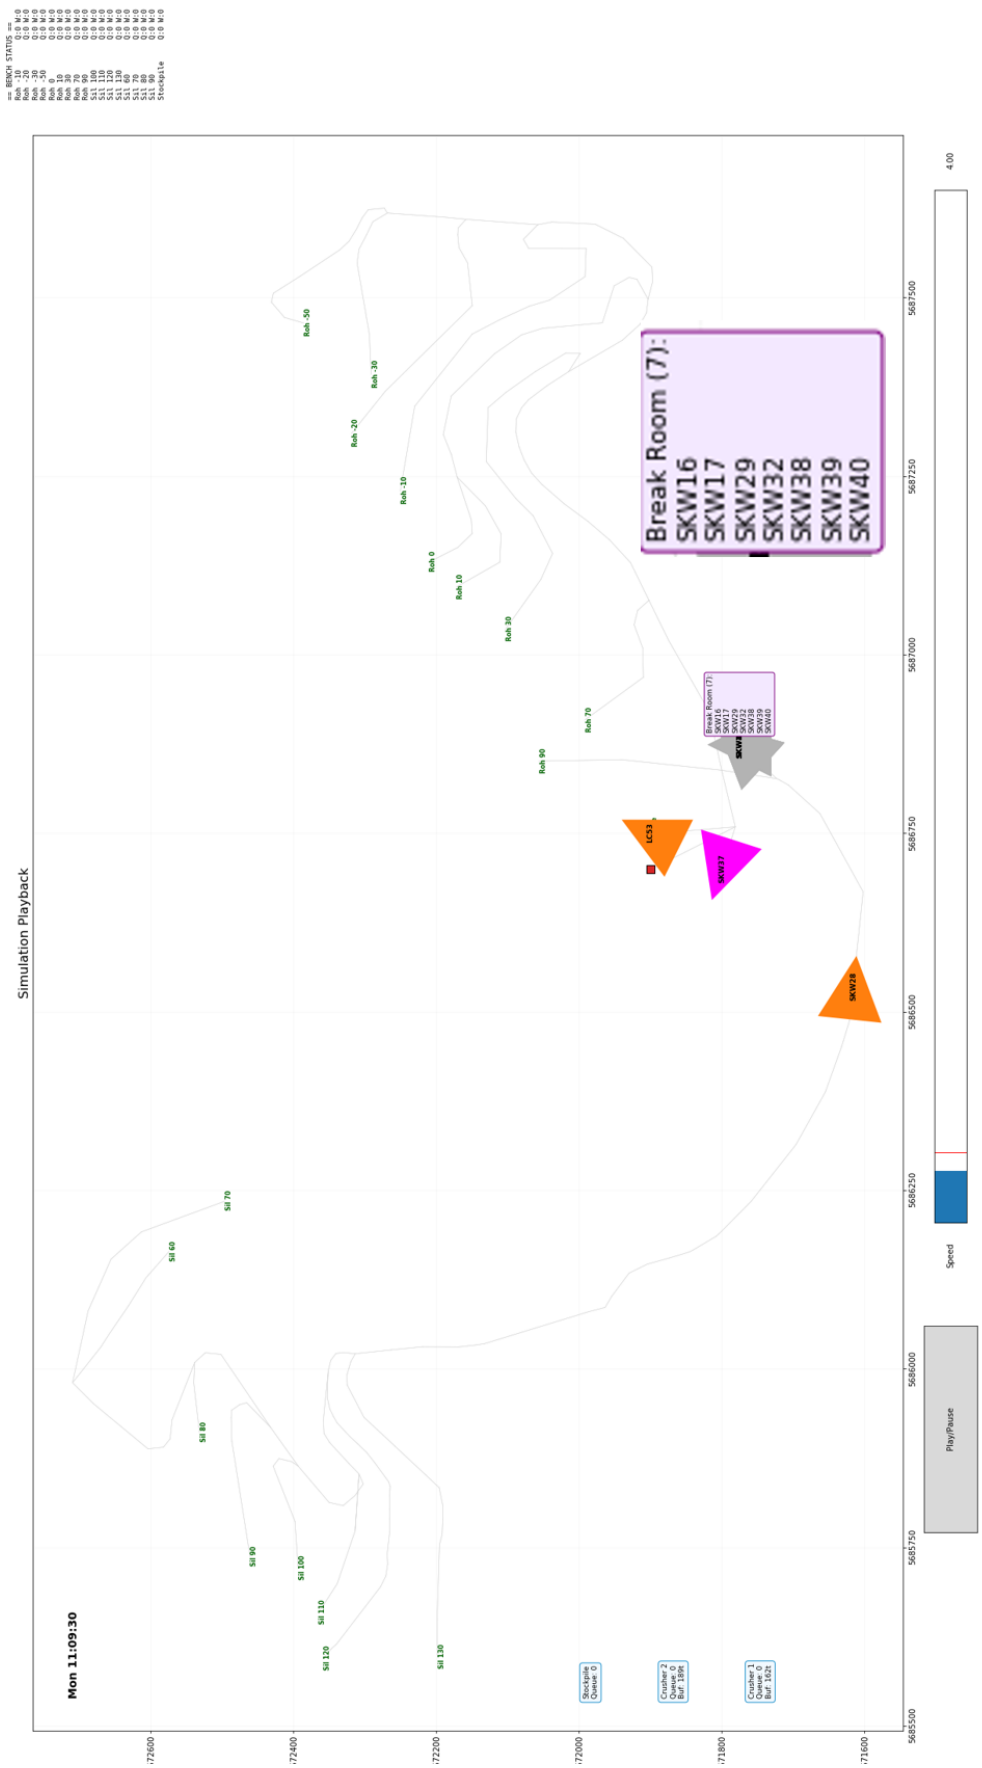

Figure G.6: Playback snapshot nine minutes into the scheduled break. Most haulers have already entered the breakroom, while one vehicle is still travelling toward the main dumping area and another has just dumped and will now move to the breakroom. This confirms that break entry is governed by the current operational state of each vehicle rather than by an instantaneous system-wide stop.

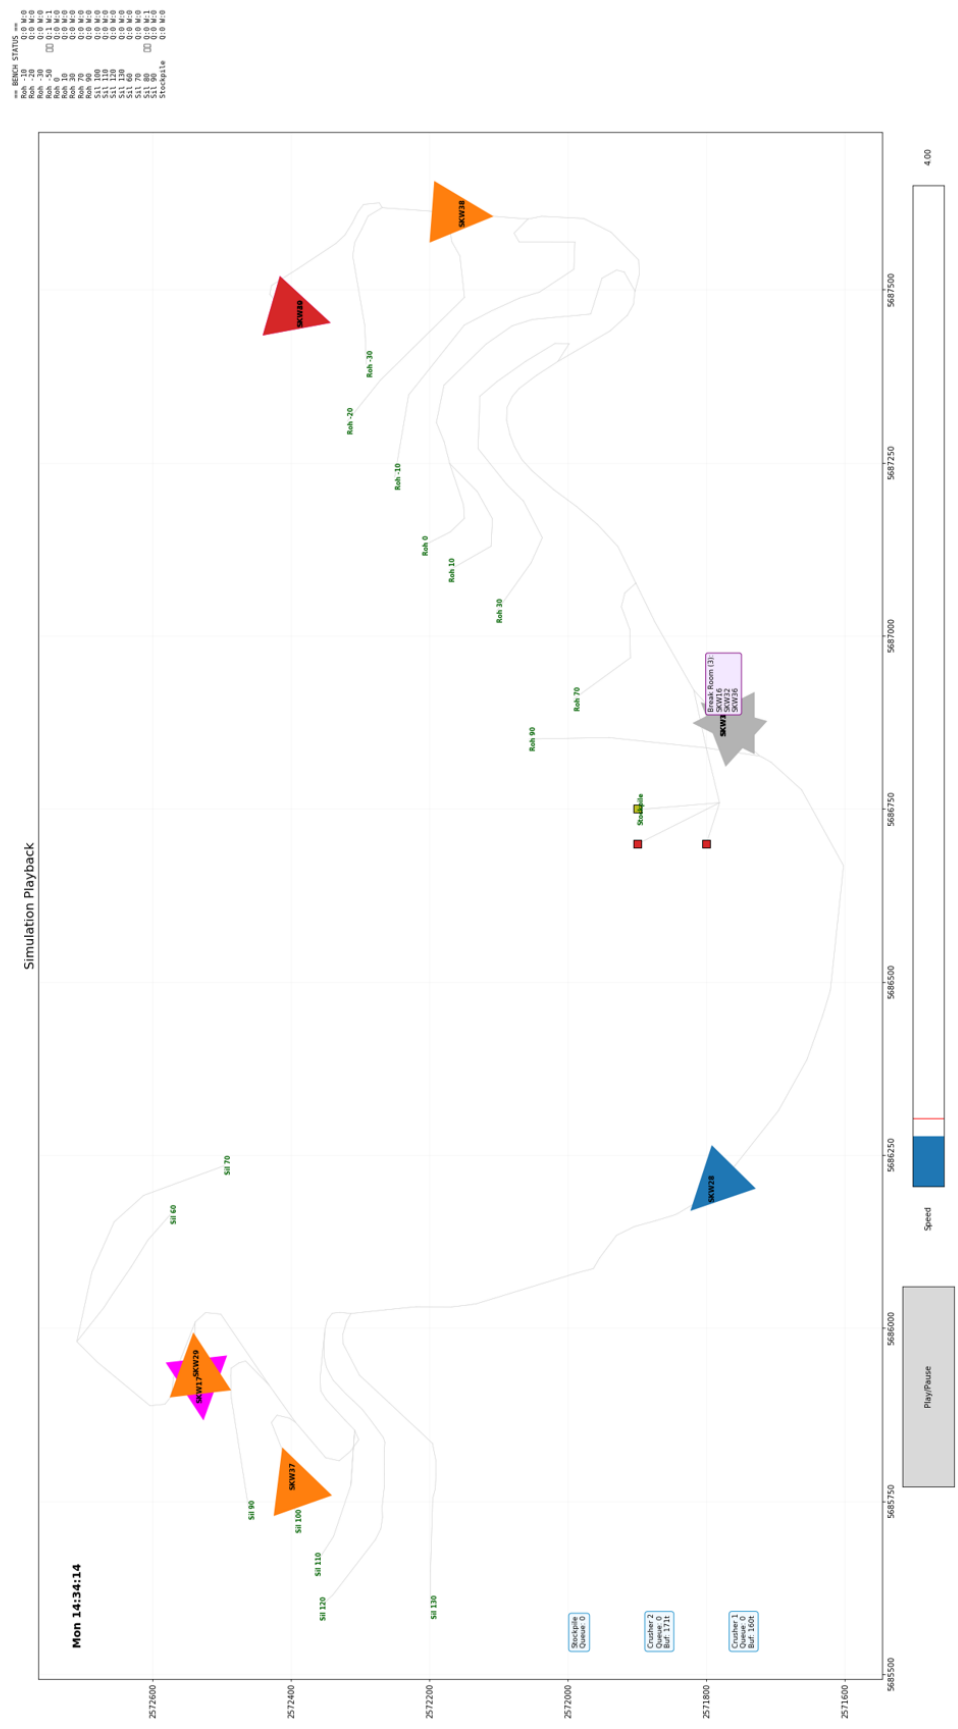

Figure G.7: Playback snapshot after the shift change. Two benches that were active previously remain active, while one previously active bench has been removed and another has been added. This confirms that the set of active benches depends on the shift schedule implemented in the DES.

H

Model Overview

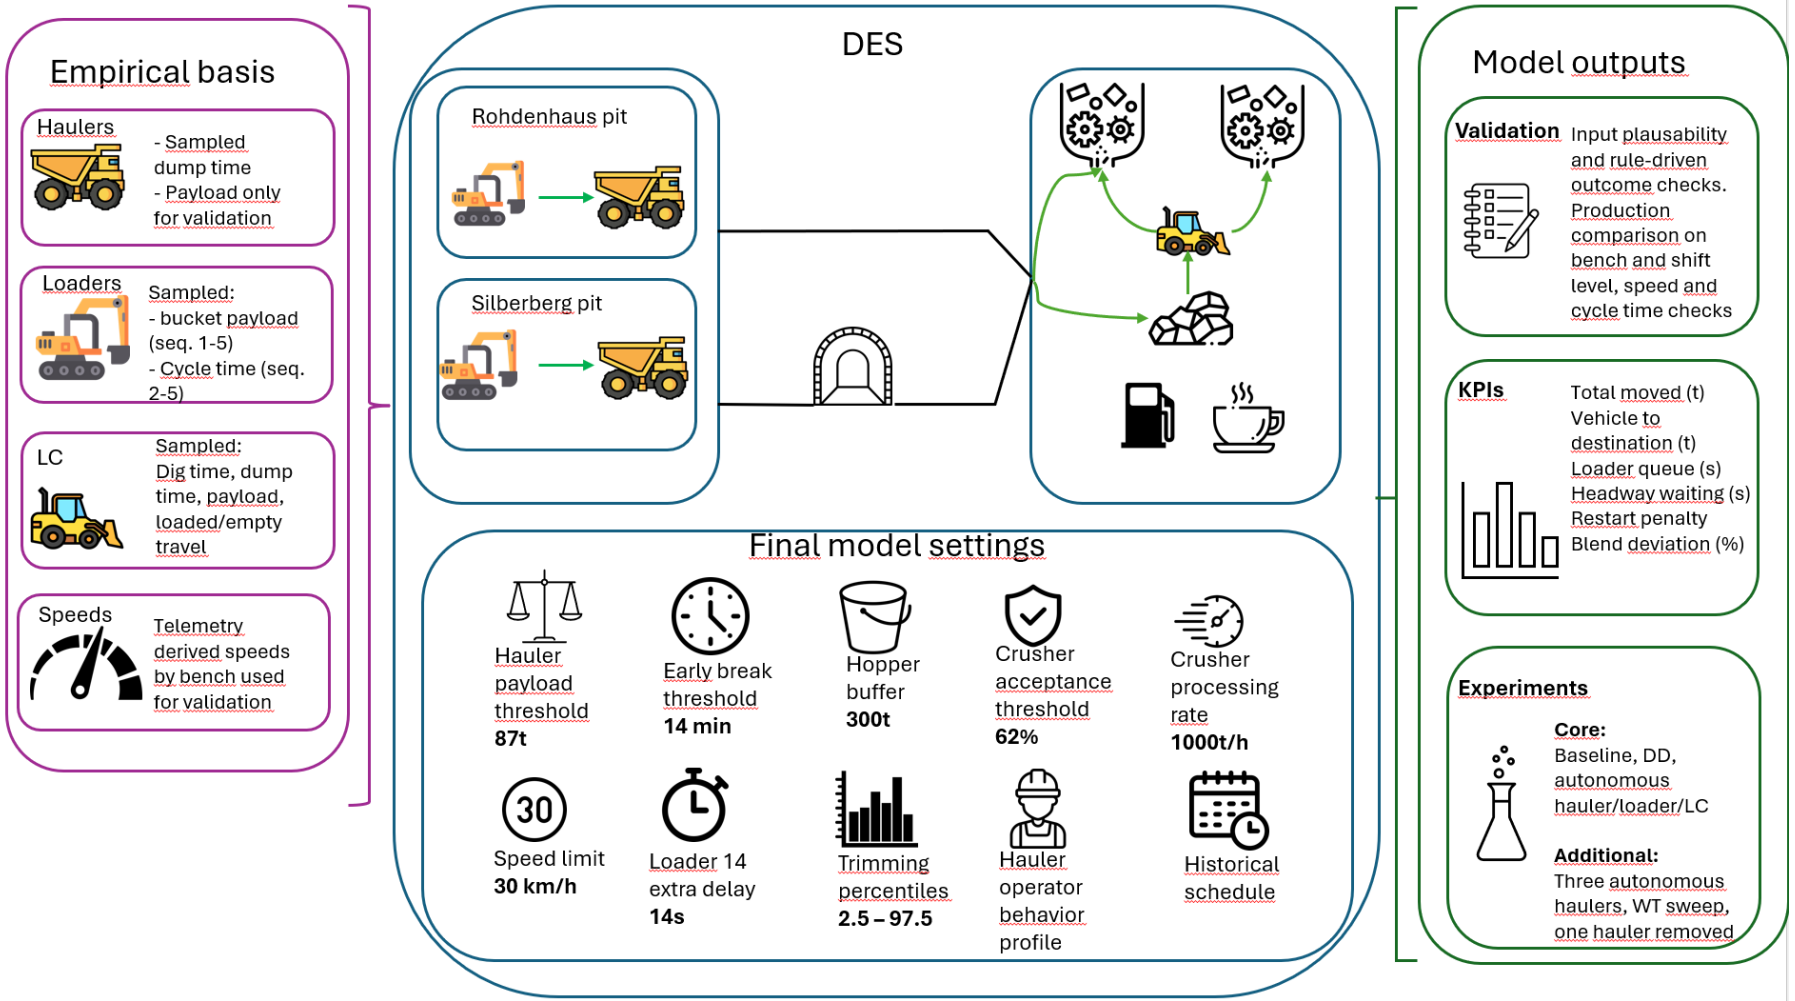

Figure H.1: Companion schematic to Figure E.1, summarizing how the DES is parameterized and what is reported in Chapter 5. The figure indicates which empirical inputs are sampled directly, which final settings define the model configuration, and how the reported outputs are organized across validation, sensitivity analysis, and scenario experiments.

I

Results

I.1. Component-level validation

I.1.1. Input sampling checks

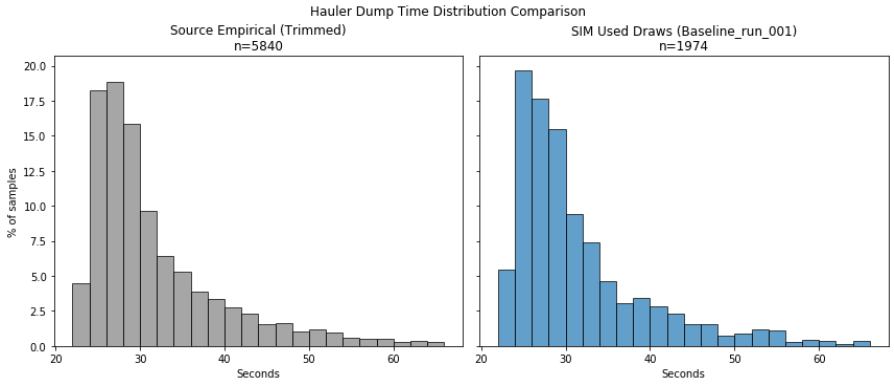

Figure I.1: Comparison between dump time input distribution and drawn values (Week 2 baseline)

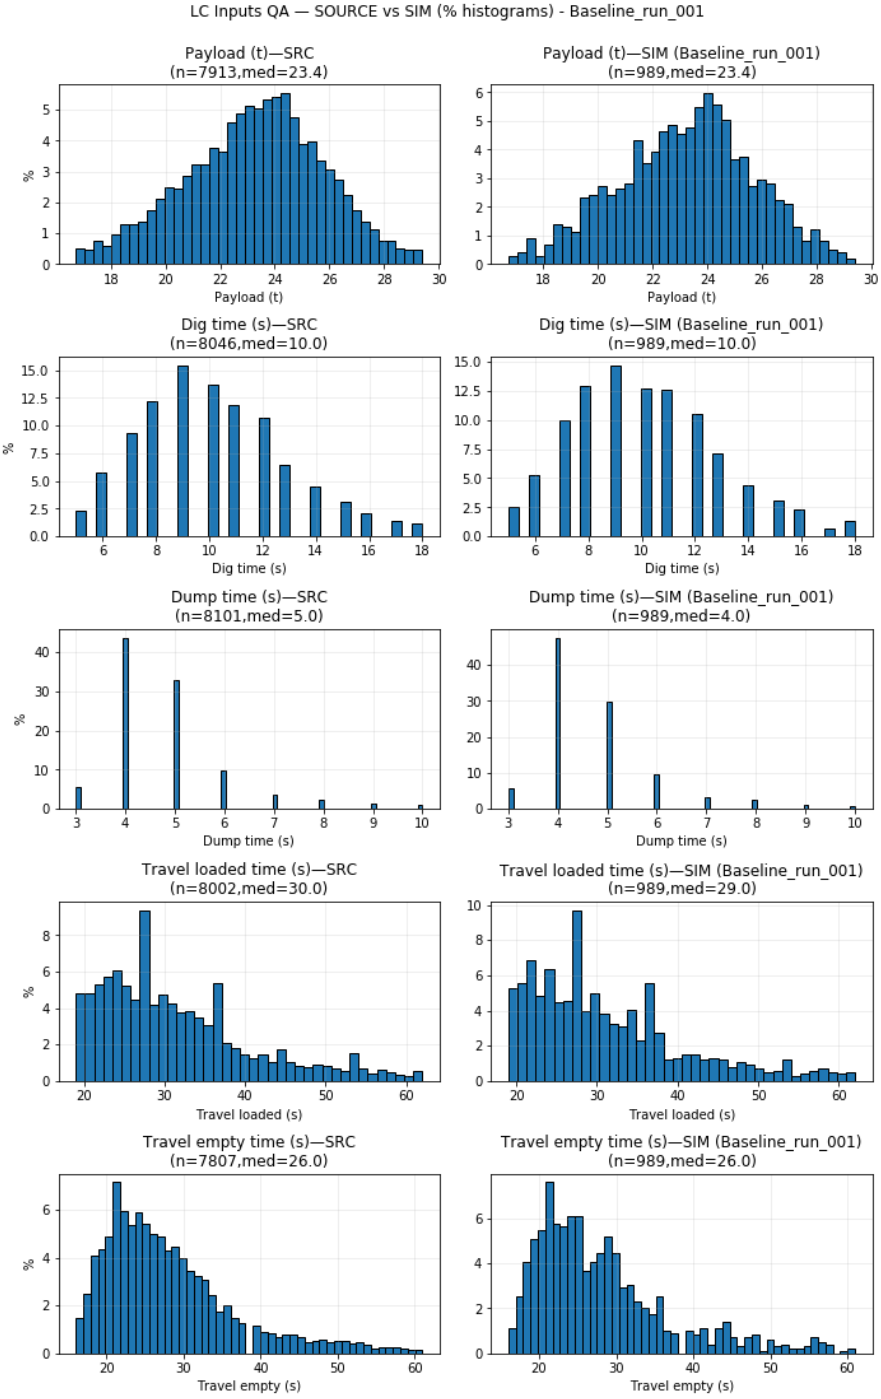

Figure I.2: Load-and-carry input distributions versus simulated draws per activity (Week 2 baseline).

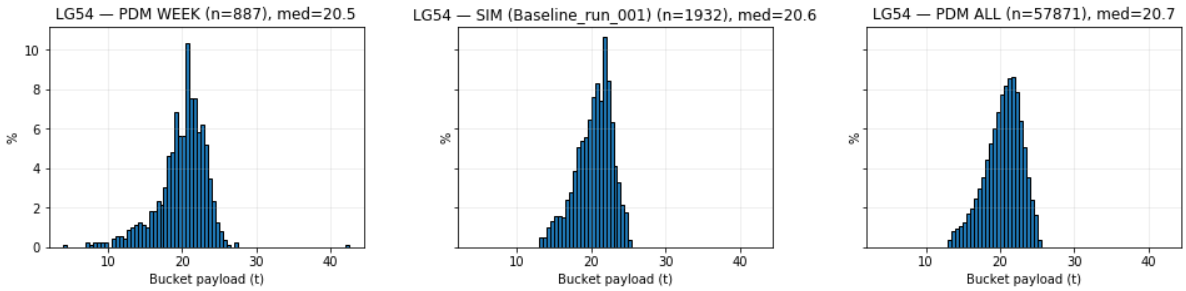

Figure I.3: Loader 54 bucket payload: empirical input versus simulated draws (Week 2 baseline).

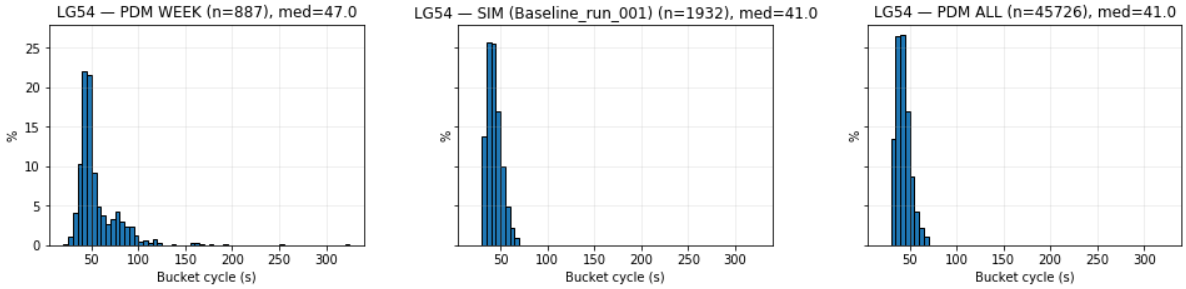

Figure I.4: Loader 54 bucket cycle time: empirical input versus simulated draws (Week 2 baseline).

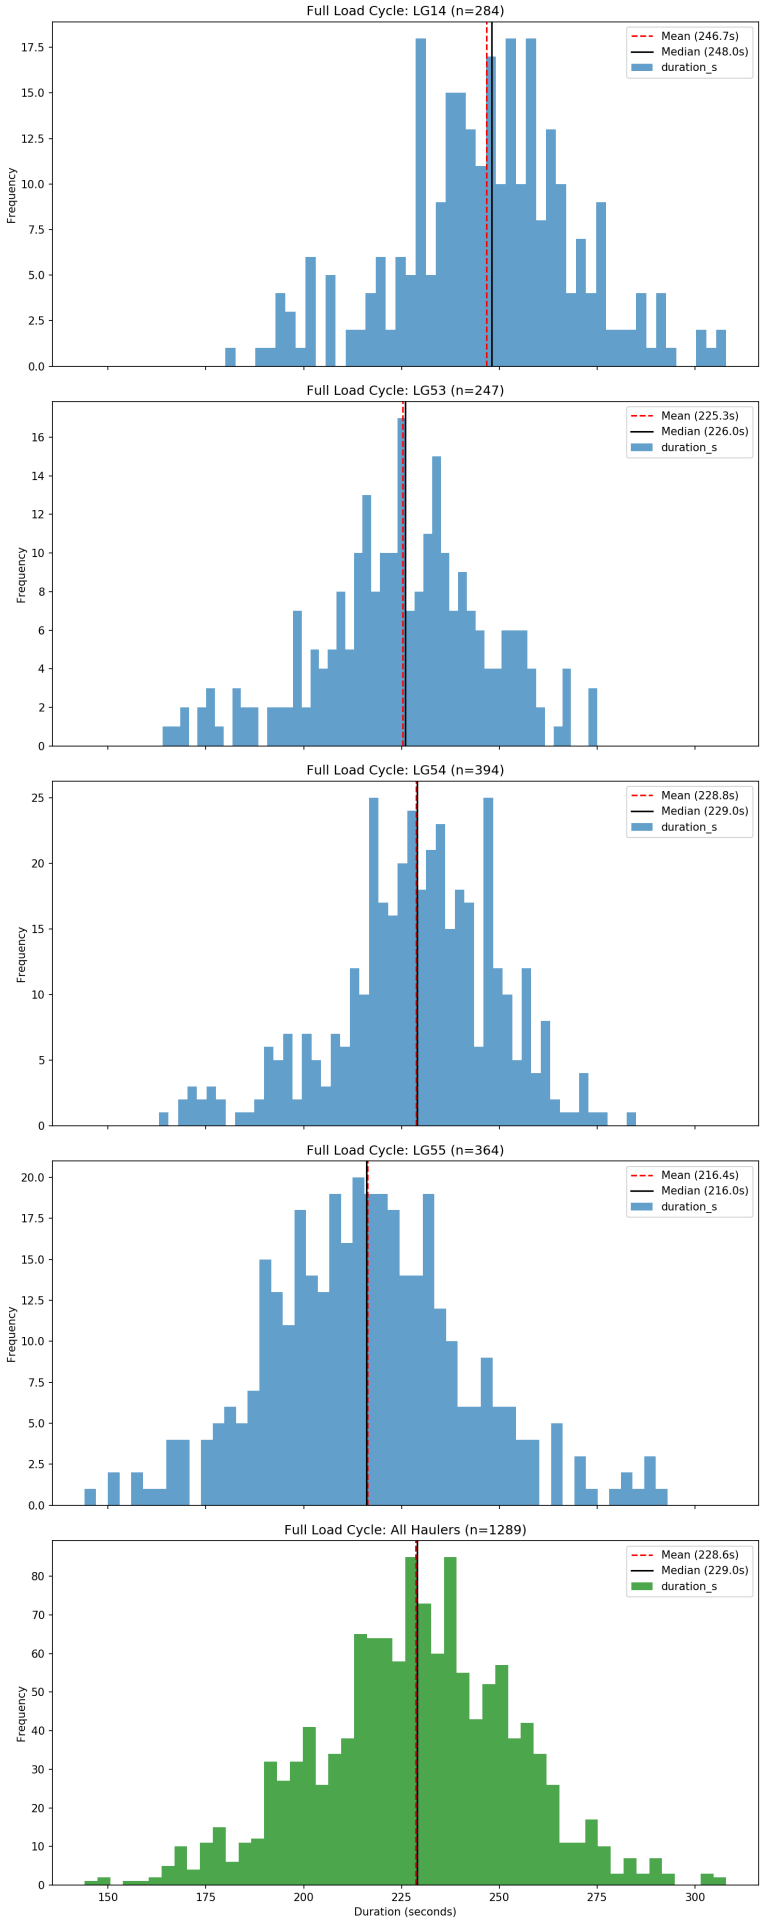

Figure I.5: Full load cycle time distributions for all modeled loaders (Week 2 baseline).

I.1.2. Rule-driven outcome checks

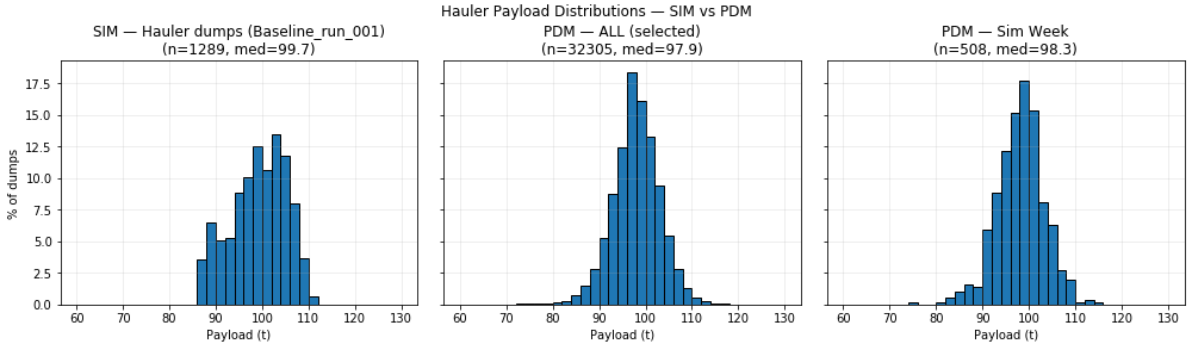

Figure I.6: Hauler payloads: telemetry reference versus simulated outcomes under bucket-by-bucket loading and the load-full stopping rule (Week 2 baseline).

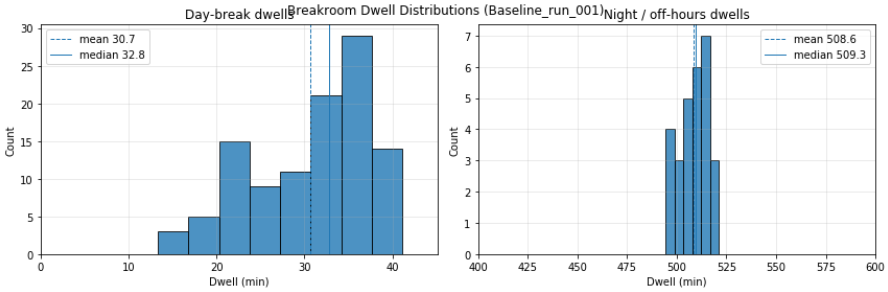

Figure I.7: Break-room dwell times during day breaks and overnight closure under the implemented break/shift-change logic (Week 2 baseline).

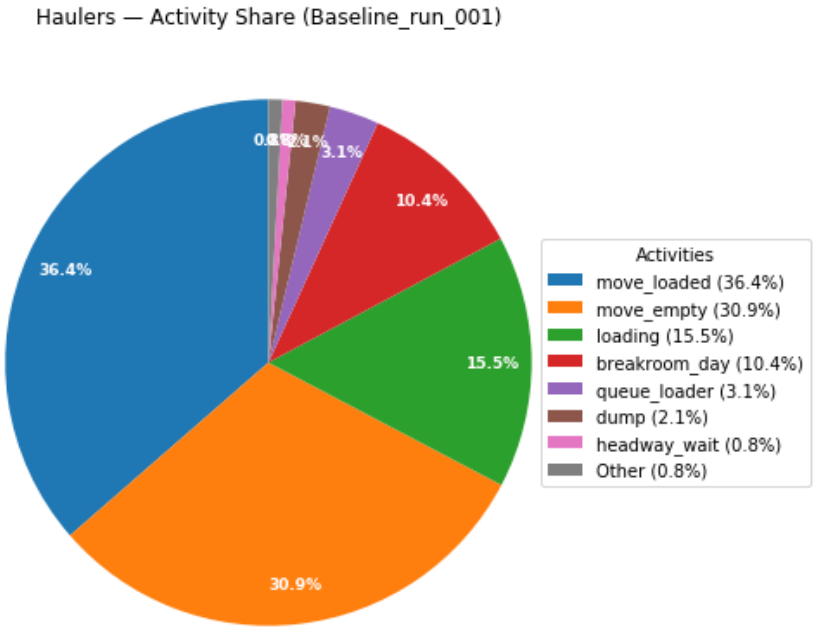

Figure I.8: Time allocation of hauler activities (Week 2 baseline).

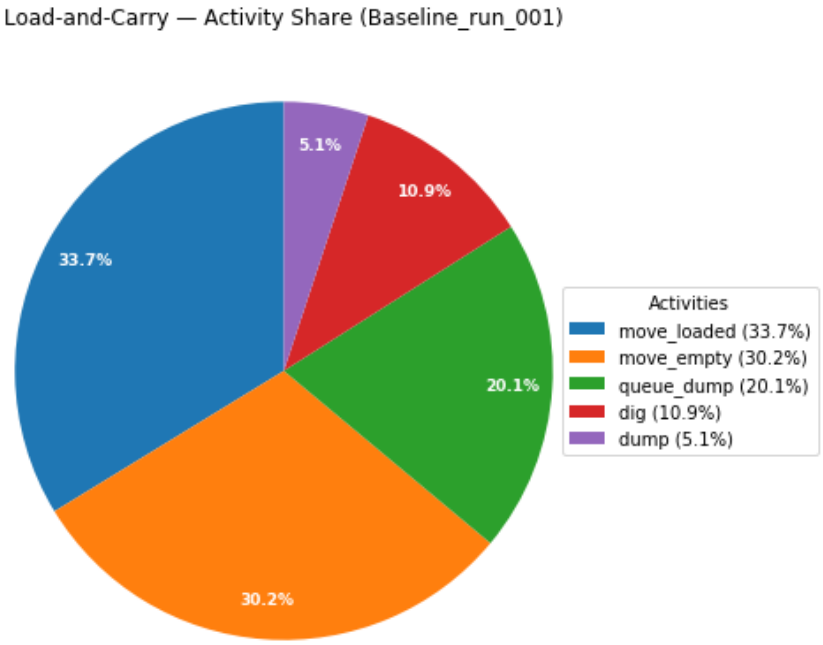

Figure I.9: Time allocation of load-and-carry activities (Week 2 baseline).

I.1.3. Qualitative behaviour traces

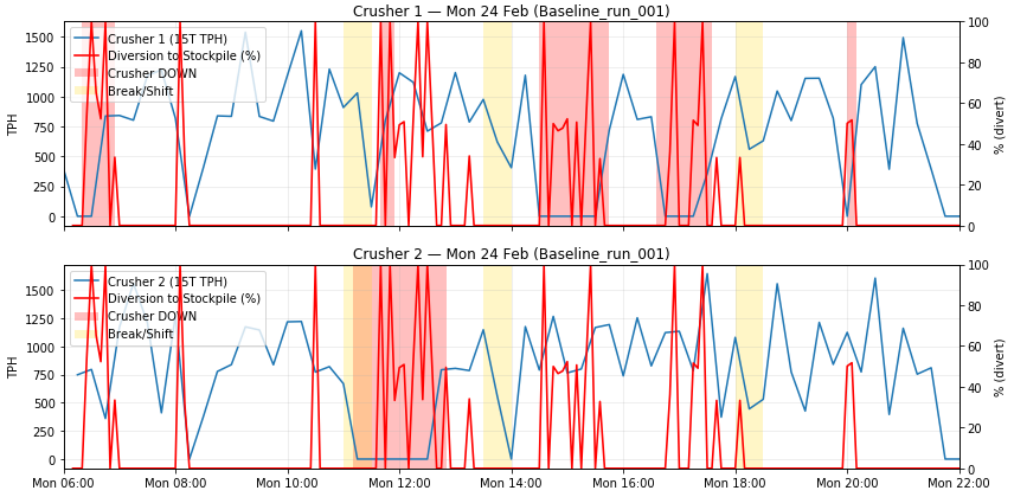

Figure I.10: Time-series behaviour during Week 2 baseline, illustrating crusher feed dynamics and diversion behaviour.

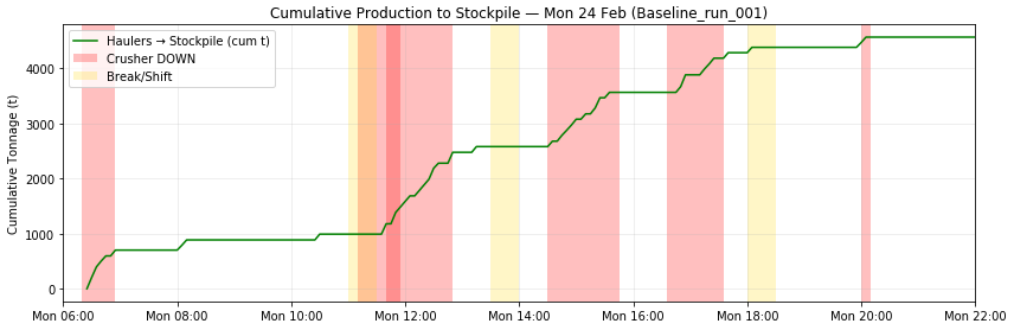

Figure I.11: Cumulative stockpiled material during Week 2 baseline.

I.2. System-level validation

I.3. Deepdive checks: speeds and cycle times

Table I.5: Loaded Haul Speed Validation: Historical vs. Weekly Simulation

Location	Real (Hist) km/h	Week 1 (Generalization) Real	Sim	Dev (%)	Week 2 (Calibration) Real	Sim	Dev (%)
Roh -10	10.7	11.6	10.5	-9.5%	11.7	10.6	-9.4%
Roh -20	10.8	11.6	10.7	-7.8%	11.4	10.8	-5.3%
Roh -50	10.8	10.5	10.3	-1.9%	–	10.4	–
Roh 90	7.1	5.9	8.7	+47.5%	6.3	7.5	+19.0%
Sil 100	14.0	14.3	13.8	-3.5%	–	–	–
Sil 110	14.3	14.4	15.0	+4.2%	14.7	14.3	-2.7%
Sil 120	14.3	–	–	–	14.5	15.6	+7.6%
Sil 60	13.7	–	–	–	14.8	12.5	-15.5%
Sil 70	13.8	13.1	12.3	-6.1%	14.5	12.5	-13.8%
Sil 80	13.7	13.7	13.5	-1.5%	–	–	–

*Speed values in km/h. Deviations calculated as (Sim - Real) / Real.

Table I.1: Week 1 production deviation (Sim – Real) by bench and shift (hauler stream).

Shift	Roh -10	Roh -20	Roh -50	Roh 90	Sil 70	Sil 80	Sil 100	Sil 110	Shift Total
Early (Mon)	–	–	+25.1% (+993)	+2.1% (+70)	+3.5% (+131)	–	–	–	**+10.8%** (+1,194)
Late (Mon)	–	–	-1.4% (-78)	–	–	-10.1% (-573)	-4.1% (-160)	–	**-5.4%** (-811)
Early (Tue)	+9.2% (+411)	–	–	–	–	+9.7% (+430)	+1.9% (+65)	–	**+7.3%** (+906)
Late (Tue)	–	–	-6.3% (-354)	–	–	+47.5% (+1,126)	+2.3% (+88)	–	**+7.3%** (+860)
Early (Wed)	–	–	-7.3% (-413)	–	–	+58.9% (+1,245)	+11.6% (+526)	–	**+11.0%** (+1,359)
Late (Wed)	–	–	+4.5% (+228)	–	–	+4.1% (+201)	+5.9% (+211)	–	**+4.7%** (+639)
Early (Thu)	–	–	+31.1% (+1,233)	–	–	-6.9% (-249)	-4.6% (-163)	–	**+7.4%** (+822)
Late (Thu)	–	-3.6% (-200)	–	–	–	+4.5% (+221)	+1.6% (+59)	–	**+0.6%** (+80)
Early (Fri)	–	-28.1% (-1,169)	–	–	–	+16.6% (+394)	+19.3% (+478)	–	**-3.3%** (-296)
Late (Fri)	-5.0% (-229)	–	–	–	+18.0% (+575)	–	–	+37.2% (+920)	**+12.4%** (+1,266)
Bench Total	**+2.0%** (+182)	**-14.1%** (-1,369)	**+5.4%** (+1,609)	**+2.1%** (+70)	**+10.2%** (+707)	**+9.2%** (+2,796)	**+3.8%** (+1,104)	**+37.2%** (+920)	**+5.0%** (+6,018)

Values format: Deviation % (Deviation tonnes). '–' indicates bench inactive in that shift.

Table I.2: Week 2 production deviation (Sim – Real) by bench and shift (hauler stream).

Shift	Roh -10	Roh -20	Roh -50	Roh 90	Sil 60	Sil 70	Sil 110	Sil 120	Shift Total
Early (Mon)	-6.4% (-308)	–	–	+57.1% (+2,375)	–	+3.5% (+94)	–	–	**+18.5%** (+2,160)
Late (Mon)	-5.1% (-280)	–	–	-6.2% (-227)	–	+4.4% (+182)	–	–	**-2.5%** (-325)
Early (Tue)	-3.7% (-188)	–	–	+44.4% (+2,021)	–	+2.4% (+63)	–	–	**+15.5%** (+1,896)
Late (Tue)	-13.4% (-781)	–	–	+32.9% (+587)	–	–	-8.5% (-550)	–	**-5.3%** (-744)
Early (Wed)	-3.3% (-165)	–	–	–	–	+3.2% (+86)	+16.3% (+532)	–	**+4.2%** (+453)
Late (Wed)	–	–	+9.4% (+455)	–	–	+4.1% (+119)	–	-2.1% (-84)	**+4.2%** (+490)
Early (Thu)	–	–	+20.6% (+877)	–	+6.5% (+179)	–	–	+4.2% (+227)	**+10.4%** (+1,283)
Late (Thu)	–	–	–	-3.1% (-215)	+7.1% (+300)	–	–	-8.9% (-388)	**-1.9%** (-303)
Early (Fri)	–	–	+19.1% (+660)	–	+35.0% (+624)	–	–	+11.6% (+320)	**+20.0%** (+1,604)
Late (Fri)	–	-0.1% (-5)	–	–	–	-5.1% (-196)	–	+10.2% (+315)	**+1.0%** (+114)
Bench Total	**-6.6%** (-1,722)	**-0.1%** (-5)	**+15.8%** (+1,993)	**+21.4%** (+4,540)	**+12.6%** (+1,103)	**+1.8%** (+347)	**-0.2%** (-18)	**+2.0%** (+389)	**+5.5%** (+6,628)

Values format: Deviation % (Deviation tonnes). '–' indicates bench inactive in that shift.

Table I.3: Week 1 shift-level production comparison: real-world records versus simulation (baseline mean).

	Hauler-to-crusher			Hauler-to-stockpile			LC-to-crusher			Total moved			Total crushed		
Shift	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**
Early (Mon)	5,839	7,776.8	+1,937.8	5,244	4,500.1	-743.9	3,783	6,295.8	+2,512.8	14,866	18,572.6	+3,706.6	9,622	14,072.4	+4,450.4
Late (Mon)	11,892	7,896.0	-3,996.0	3,186	6,371.2	+3,185.2	2,897	5,351.0	+2,454.0	17,975	19,618.0	+1,643.0	14,789	13,247.1	-1,541.9
Early (Tue)	6,649	9,444.7	+2,795.7	5,705	3,815.2	-1,889.8	4,228	3,434.3	-793.7	16,582	16,694.1	+112.1	10,877	12,878.9	+2,001.9
Late (Tue)	7,385	7,277.1	-107.9	4,450	5,417.7	+967.7	4,070	5,337.7	+1,267.7	15,905	18,032.7	+2,127.7	11,455	12,615.0	+1,160.0
Early (Wed)	10,626	10,337.7	-288.3	4,155	3,332.5	-822.5	3,300	2,129.7	-1,170.3	18,081	15,799.7	-2,281.3	13,926	12,467.3	-1,458.7
Late (Wed)	9,298	8,030.3	-1,267.7	4,236	6,142.8	+1,906.8	2,363	5,418.8	+3,055.8	15,897	19,591.9	+3,694.9	11,661	13,449.1	+1,788.1
Early (Thu)	7,310	10,962.8	+3,652.8	3,840	1,008.9	-2,831.1	2,058	0.0	-2,058.0	13,208	11,971.7	-1,236.3	9,368	10,962.8	+1,594.8
Late (Thu)	12,956	12,591.2	-364.8	1,263	1,707.8	+444.8	0	0.0	+0.0	14,219	14,299.0	+80.0	12,956	12,591.2	-364.8
Early (Fri)	7,230	7,492.1	+262.1	1,774	1,215.9	-558.1	0	0.0	+0.0	9,004	8,708.0	-296.0	7,230	7,492.1	+262.1
Late (Fri)	7,171	10,260.2	+3,089.2	3,045	1,221.6	-1,823.4	0	0.0	+0.0	10,216	11,481.8	+1,265.8	7,171	10,260.2	+3,089.2

Diff = Sim – Real (tonnes).

Table I.4: Week 2 shift-level production comparison: real-world records versus simulation (baseline mean).

	Hauler-to-crusher			Hauler-to-stockpile			LC-to-crusher			Total moved			Total crushed		
Shift	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**	**Real**	**Sim**	**Diff**
Early (Mon)	10,701	12,202.6	+1,501.6	2,453	1,626.7	-826.3	1,273	1,074.4	-198.6	14,427	14,903.5	+476.5	11,974	13,276.8	+1,302.8
Late (Mon)	9,719	11,391.1	+1,672.1	1,950	1,527.6	-422.4	1,067	1,045.5	-21.5	12,736	13,964.2	+1,228.2	10,786	12,436.7	+1,650.7
Early (Tue)	8,373	11,234.9	+2,861.9	3,890	2,923.9	-966.1	1,514	1,833.3	+319.3	13,777	15,992.3	+2,215.3	9,887	13,068.3	+3,181.3
Late (Tue)	10,249	9,970.1	-278.9	3,831	3,365.6	-465.4	2,750	2,627.8	-122.2	16,830	15,963.7	-866.3	12,999	12,598.1	-400.9
Early (Wed)	4,874	7,752.8	+2,878.8	6,024	3,568.8	-2,455.2	5,438	5,589.8	+151.8	16,307	16,911.4	+604.4	10,283	13,342.7	+3,059.7
Late (Wed)	6,287	7,295.7	+1,008.7	5,473	4,984.6	-488.4	5,473	6,096.1	+623.1	17,233	18,376.5	+1,143.5	11,760	13,391.8	+1,631.8
Early (Thu)	11,730	11,818.6	+88.6	624	1,818.4	+1,194.4	228	496.1	+268.1	12,582	14,133.3	+1,551.3	11,958	12,314.9	+356.9
Late (Thu)	14,152	12,673.2	-1,478.8	1,461	2,636.8	+1,175.8	0	10.1	+10.1	15,613	15,320.1	-292.9	14,152	12,683.6	-1,468.4
Early (Fri)	11,380	7,131.3	-4,248.7	198	2,474.0	+2,276.0	0	3,159.4	+3,159.4	11,578	12,764.4	+1,186.4	11,380	10,290.4	-1,089.6
Late (Fri)	4,847	9,735.3	+4,888.3	2,974	1,956.2	-1,017.8	0	1,288.3	+1,288.3	10,597	12,979.9	+2,382.9	7,623	11,023.7	+3,400.7

Diff = Sim – Real (tonnes).

Table I.6: Empty Haul Speed Validation: Historical vs. Weekly Simulation

Location	Real (Hist) km/h	Week 1 (Generalization) Real	Sim	Dev (%)	Week 2 (Calibration) Real	Sim	Dev (%)
Roh -10	14.4	14.8	14.5	-2.3%	15.5	14.6	-5.9%
Roh -20	14.7	14.3	15.6	+8.7%	15.2	15.7	+3.1%
Roh -50	14.8	14.4	16.5	+14.6%	–	16.6	–
Roh 90	8.2	3.5	8.5	+140.8%	6.3	7.5	+18.4%
Sil 100	16.9	17.3	15.9	-8.5%	–	–	–
Sil 110	16.8	18.4	16.8	-8.7%	14.7	14.3	-3.2%
Sil 120	16.2	–	–	–	14.5	15.6	+7.6%
Sil 60	16.9	17.4	16.8	-3.2%	14.8	12.5	-15.3%
Sil 70	16.6	15.5	16.0	+2.7%	14.5	12.5	-13.5%
Sil 80	16.7	16.8	16.4	-2.0%	–	–	–

*Speed values in km/h. Deviations calculated as (Sim - Real) / Real. '–' indicates insufficient data.

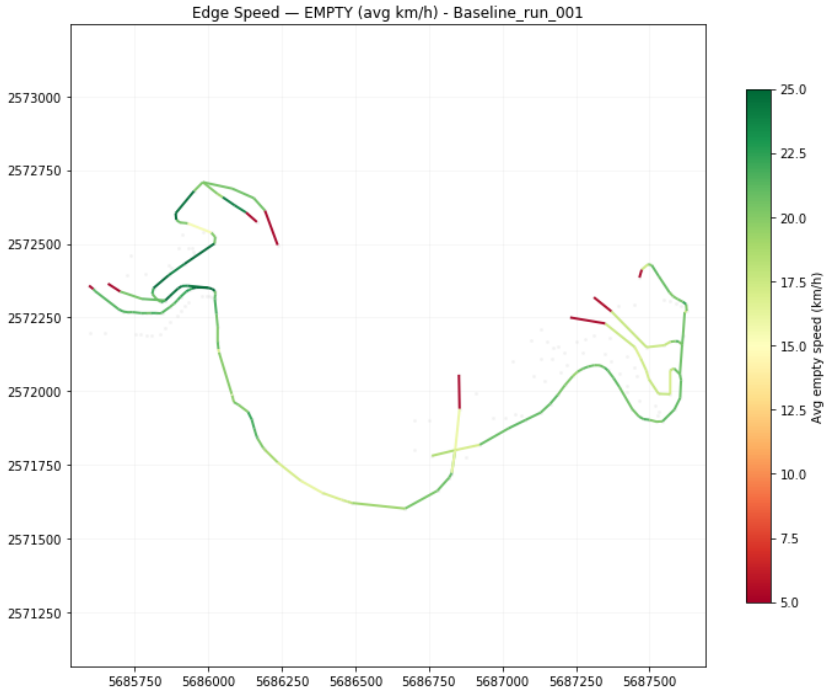

Figure I.12: Stamped Empty Speeds

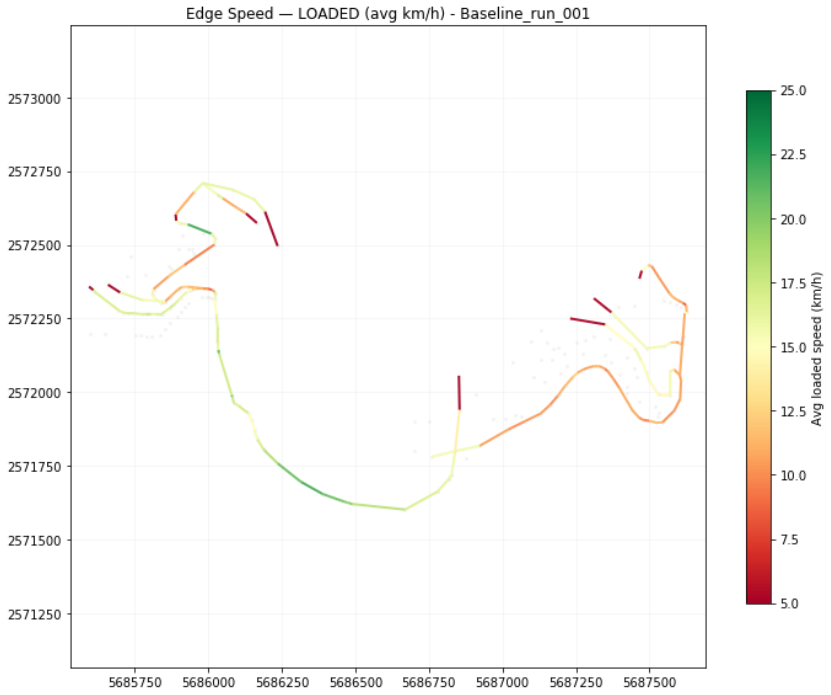

Figure I.13: Stamped Loaded Speed

I.4. Sensitivity Analysis

Table I.7: Week 1 sensitivity analysis: production outcomes relative to the Baseline configuration (one-factor-at-a-time).

Setting	Mean Tonnage	Min	Max	Gain vs Baseline (%)
HT_95	159,921.9	158,544	160,992	+3.33
HT_80	145,038.9	142,741	146,322	-6.29
CR_TPH_1200	157,223.4	155,639	158,234	+1.59
CR_TPH_800	152,037.3	149,509	153,567	-1.77
BUF_200	158,256.9	155,545	159,694	+2.25
BUF_400	155,111.5	153,211	155,923	+0.22
HBAYS_2	155,413.7	153,087	156,942	+0.42
DT_single_block_daily	155,788.5	154,399	157,187	+0.66
DT_deterministic_90_10	154,118.1	152,336	155,776	-0.42
SPEED_MULT_1p1	160,870.8	158,720	162,763	+3.94
SPEED_MULT_0p9	148,884.0	146,892	150,177	-3.80
SPEED_B_MULT_1	158,483.5	157,352	160,854	+2.40
SPEED_B_MULT_5	147,734.9	146,416	148,581	-4.55
SPD_KAPPA_0p05	157,986.0	156,659	158,881	+2.08
SPD_KAPPA_0p25	153,584.8	151,712	155,092	-0.77
ACC_0p7	156,381.3	154,243	157,680	+1.04
ACC_0p55	154,602.3	152,425	156,240	-0.11
SIG_LC53_2	166,419.7	165,374	167,948	+7.53
SIG_LC53_3	166,312.4	163,637	168,376	+7.46
SIG_LC53_5	166,697.8	165,224	168,624	+7.71
SIG_LC53_6	166,357.3	164,215	167,575	+7.49
SIG_SKW36_2	156,602.2	155,049	157,943	+1.18
SIG_SKW36_3	156,337.3	154,395	157,922	+1.01
SIG_SKW36_5	156,532.7	154,152	157,877	+1.14
SIG_SKW36_6	156,361.2	154,720	157,979	+1.03
RANGL_10	155,765.4	154,613	157,470	+0.64
RANGL_30	153,986.9	152,819	155,563	-0.51
RANGH_5	155,503.9	153,311	156,727	+0.47
RANGH_15	154,537.2	153,098	156,140	-0.15
LG14EX_0	154,960.1	153,370	156,028	+0.12
LG14EX_40	154,100.2	153,310	154,803	-0.43
SIG_LG54_2	155,612.1	154,178	157,576	+0.54
SIG_LG54_3	155,867.0	154,566	157,433	+0.71
SIG_LG54_5	155,356.1	154,340	156,928	+0.38
SIG_LG54_6	155,577.9	153,419	157,987	+0.52
TRIM_5p0_95	155,760.3	154,142	156,715	+0.64
TRIM_1_99	154,023.0	152,742	155,431	-0.48
TRIM_0_100	149,898.5	138,723	154,058	-3.15
Baseline	154,769.9	152,819	156,623	0.00

Baseline is shown for reference. Rows identical to the Baseline configuration are omitted.

Table I.8: Week 2 sensitivity analysis: production outcomes relative to the Baseline configuration (one-factor-at-a-time).

Setting	Mean Tonnage	Min	Max	Gain vs Baseline (%)
HT_95	156,052.8	153,555	157,864	+3.14
HT_80	141,586.0	140,274	143,307	-6.43
CR_TPH_1200	153,219.6	151,578	154,508	+1.26
CR_TPH_800	149,137.7	147,691	150,409	-1.44
BUF_200	153,031.6	150,898	154,169	+1.14
BUF_400	151,757.8	150,249	153,509	+0.30
HBAYS_2	152,633.5	151,259	154,155	+0.88
DT_single_block_daily	151,931.3	150,490	153,447	+0.41
DT_deterministic_90_10	150,810.0	150,150	151,929	-0.33
SPEED_MULT_1p1	156,919.9	155,776	158,518	+3.71
SPEED_MULT_0p9	145,173.0	143,646	146,519	-4.06
SPEED_BMULT_1	154,134.4	152,353	155,377	+1.87
SPEED_BMULT_5	146,011.1	145,110	147,504	-3.50
SPD_KAPPA_0p05	154,059.4	153,129	154,837	+1.82
SPD_KAPPA_0p25	148,691.9	147,770	150,225	-1.73
ACC_0p7	152,235.4	149,814	153,314	+0.61
ACC_0p55	151,790.0	151,071	152,646	+0.32
SIG_LC53_2	157,905.9	156,577	158,950	+4.36
SIG_LC53_3	158,271.4	157,131	159,113	+4.60
SIG_LC53_5	158,054.1	156,758	159,877	+4.46
SIG_LC53_6	157,876.4	156,624	158,878	+4.34
SIG_SKW36_2	148,400.8	147,320	149,926	-1.92
SIG_SKW36_3	148,430.8	145,975	149,718	-1.90
SIG_SKW36_5	148,563.6	147,217	149,891	-1.81
SIG_SKW36_6	148,294.5	147,028	149,897	-1.99
RANGL_10	152,582.4	150,790	153,875	+0.84
RANGL_30	150,566.1	148,802	151,880	-0.49
RANGH_5	152,169.5	150,255	153,422	+0.57
RANGH_15	150,961.1	149,429	152,794	-0.23
LG14EX_0	151,934.9	151,052	153,272	+0.41
LG14EX_40	151,043.4	150,053	151,869	-0.18
SIG_LG54_2	152,073.7	150,706	153,839	+0.51
SIG_LG54_3	151,928.8	149,403	153,387	+0.41
SIG_LG54_5	152,150.3	150,272	153,270	+0.56
SIG_LG54_6	152,142.5	150,803	153,406	+0.55
TRIM_5p0_95	152,314.4	150,008	154,184	+0.66
TRIM_1_99	150,683.9	148,689	152,006	-0.41
TRIM_0_100	149,870.6	147,820	151,293	-0.95
Baseline	151,309.0	149,421	152,487	0.00

Baseline is shown for reference.

Table I.9: Week 1 sensitivity analysis: key waiting-time KPIs (seconds).

Setting	LQ pos	LQ incl0	SQ pos	SQ incl0	CQ-LC incl0	Headway pos	Restart
HT_95	148.88	52.68	24.10	2.73	24.95	18.77	1.40
HT_80	135.96	45.98	25.57	3.22	21.14	16.92	1.28
CR_TPH_1200	145.17	50.80	23.89	2.52	16.59	18.06	1.36
CR_TPH_800	140.97	50.57	23.24	3.11	34.54	17.22	1.37
BUF_200	142.17	50.20	24.49	2.93	13.20	17.65	1.37
BUF_400	145.49	49.79	25.39	3.05	22.34	17.18	1.35
HBAYS_2	143.52	51.69	25.02	2.64	20.12	16.78	1.36
DT_single_block_daily	144.58	49.96	23.92	2.76	21.90	17.33	1.36
DT_deterministic_90_10	142.67	49.44	25.50	3.12	27.33	17.25	1.35
SPEED_MULT_1p1	140.27	49.83	23.26	2.68	26.08	18.19	1.38
SPEED_MULT_0p9	148.96	48.47	25.18	2.64	21.33	17.33	1.33
SPEED_B_MULT_1	139.69	49.25	24.52	2.82	25.00	17.61	1.38
SPEED_B_MULT_5	148.20	48.71	23.88	2.81	21.83	18.20	1.37
SPD_KAPPA_0p05	97.87	25.60	23.31	2.64	24.72	32.69	2.59
SPD_KAPPA_0p25	157.63	54.48	25.17	3.52	23.46	14.10	1.17
ACC_0p7	146.88	50.56	23.72	2.86	18.21	18.08	1.36
ACC_0p55	142.55	49.63	24.28	2.76	24.40	17.55	1.34
SIG_LC53_2	143.58	50.65	24.81	2.78	22.99	17.98	1.34
SIG_LC53_3	142.76	49.69	24.56	2.84	23.20	17.82	1.37
SIG_LC53_5	144.14	50.66	25.37	3.09	22.79	17.50	1.37
SIG_LC53_6	142.53	50.42	23.69	2.85	23.65	17.84	1.36
SIG_SKW36_2	141.21	45.80	25.01	2.85	24.00	16.53	1.42
SIG_SKW36_3	140.20	45.41	23.96	2.62	24.33	16.59	1.41
SIG_SKW36_5	139.40	46.07	23.48	2.80	24.25	16.73	1.43
SIG_SKW36_6	140.41	45.51	25.20	2.59	24.95	16.73	1.44
RANGL_10	139.00	48.30	24.29	2.94	24.92	17.08	1.33
RANGL_30	148.29	52.18	22.39	2.47	24.13	18.75	1.37
RANGH_5	143.75	49.80	22.59	2.57	22.96	17.72	1.33
RANGH_15	143.49	48.94	26.58	3.51	24.53	17.08	1.37
LG14EX_0	140.43	48.70	23.96	2.80	24.78	17.45	1.35
LG14EX_40	145.62	51.79	25.04	2.82	24.13	17.47	1.42
SIG_LG54_2	143.89	51.30	24.02	2.85	24.61	18.03	1.35
SIG_LG54_3	142.81	48.99	23.45	2.75	24.90	17.71	1.36
SIG_LG54_5	145.07	52.04	24.70	2.92	24.14	17.92	1.36
SIG_LG54_6	142.19	51.11	23.57	2.70	24.05	17.94	1.36
TRIM_5p0_95	143.60	50.16	23.26	2.71	24.67	16.92	1.34
TRIM_1_99	145.74	50.44	25.06	3.02	24.10	18.21	1.40
TRIM_0_100	146.11	49.15	1397.24	220.59	24.42	17.69	1.41
Baseline	143.36	50.66	24.01	2.63	24.06	18.08	1.35

LQ = loader queue, SQ = stockpile queue, CQ-LC = crusher queue for load-and-carry. Baseline is shown for reference.

Table I.10: Week 2 sensitivity analysis: key waiting-time KPIs (seconds).

Setting	LQ pos	LQ incl0	SQ pos	SQ incl0	CQ-LC incl0	Headway pos	Restart
HT_95	139.17	46.23	23.63	2.46	21.29	23.21	1.46
HT_80	126.72	38.17	23.88	2.55	15.86	21.46	1.36
CR_TPH_1200	131.05	41.70	22.37	2.59	12.88	23.12	1.43
CR_TPH_800	133.51	43.18	22.84	2.47	27.35	23.10	1.40
BUF_200	132.58	43.12	22.98	2.35	11.96	23.17	1.45
BUF_400	130.63	42.38	22.92	2.23	16.92	24.26	1.39
HBAYS_2	131.50	42.65	24.74	2.11	14.69	22.90	1.48
DT_single_block_daily	134.10	43.73	25.96	2.43	15.41	23.93	1.42
DT_deterministic_90_10	133.21	43.03	22.80	2.33	19.59	23.04	1.41
SPEED_MULT_1p1	129.86	43.68	23.83	2.33	19.12	24.07	1.43
SPEED_MULT_0p9	138.09	42.20	23.70	2.29	16.75	22.57	1.43
SPEED_BMULT_1	132.46	43.06	23.21	2.32	18.78	23.81	1.43
SPEED_BMULT_5	137.04	40.94	23.59	2.27	15.81	21.51	1.40
SPD_KAPPA_0p05	95.58	24.09	22.17	1.99	19.81	39.98	2.39
SPD_KAPPA_0p25	146.84	48.10	22.67	2.58	18.42	16.41	1.19
ACC_0p7	130.83	43.11	23.28	2.25	14.78	23.46	1.42
ACC_0p55	132.97	42.80	22.88	2.29	17.82	22.60	1.42
SIG_LC53_2	135.73	42.73	25.33	2.42	16.43	22.93	1.42
SIG_LC53_3	135.87	42.10	25.13	2.29	16.25	22.08	1.46
SIG_LC53_5	136.09	41.26	25.10	2.61	16.50	22.15	1.42
SIG_LC53_6	135.65	42.34	24.18	2.55	16.48	22.79	1.42
SIG_SKW36_2	133.04	37.19	22.81	2.51	16.73	19.30	1.49
SIG_SKW36_3	133.73	39.07	24.12	2.28	17.59	20.06	1.49
SIG_SKW36_5	134.47	37.24	22.24	2.06	16.69	19.14	1.51
SIG_SKW36_6	133.16	35.81	23.19	2.38	16.75	19.13	1.54
RANGL_10	127.79	39.72	23.61	2.23	18.35	24.11	1.47
RANGL_30	137.91	44.81	23.20	2.42	17.21	23.54	1.47
RANGH_5	132.10	42.20	20.49	1.78	18.18	22.61	1.45
RANGH_15	132.66	41.62	27.57	3.30	18.60	22.21	1.44
LG14EX_0	130.70	41.69	24.96	2.45	17.99	22.74	1.42
LG14EX_40	133.91	43.33	23.30	2.44	18.34	24.49	1.50
SIG_LG54_2	134.35	44.36	22.95	2.13	18.57	23.44	1.44
SIG_LG54_3	132.62	43.12	23.17	2.20	18.91	24.93	1.46
SIG_LG54_5	137.41	43.26	23.11	2.39	18.29	23.18	1.40
SIG_LG54_6	133.48	43.96	24.59	2.36	18.92	24.70	1.44
TRIM_5p0_95	132.31	43.83	21.98	2.17	18.42	23.62	1.43
TRIM_1_99	136.16	43.53	24.55	2.40	17.97	22.20	1.43
TRIM_0_100	135.80	42.78	26.23	2.81	18.46	22.98	1.46
Baseline	133.46	44.06	20.93	2.02	17.80	23.47	1.41

LQ = loader queue, SQ = stockpile queue, CQ-LC = crusher queue for load-and-carry. Baseline is shown for reference.

I.5. Core scenarios

Table I.11: Week 1 Dynamic Dispatch compliance by shift and bench (DynamicDispatch_v1), including tonnage and blend compliance.

Shift	Bench	Target (t)	Sim. (t)	Delta (t)	Abs Delta (t)	Delta (%)	Target Mix (%)	Sim. Mix (%)	Blend Err. (pp)
Early (Mon)	Roh 90	3691	3674.8	-16.2	16.2	-0.4	29.1	30.0	0.9
Early (Mon)	Sil 70	3886	3717.4	-168.6	168.6	-4.3	30.7	30.3	0.3
Early (Mon)	Roh -50	5091	4860.9	-230.1	230.1	-4.5	40.2	39.7	0.5
Late (Mon)	Roh -50	5477	5441.7	-35.3	35.3	-0.6	38.5	38.2	0.2
Late (Mon)	Sil 80	5054	5018.6	-35.4	35.4	-0.7	35.5	35.2	0.2
Late (Mon)	Sil 100	3709	3777.4	68.4	68.4	1.8	26.0	26.5	0.5
Early (Tue)	Roh -10	4835	4768.1	-66.9	66.9	-1.4	36.7	36.5	0.2
Early (Tue)	Sil 80	4715	4632.6	-82.4	82.4	-1.7	35.7	35.5	0.3
Early (Tue)	Sil 100	3640	3666.3	26.3	26.3	0.7	27.6	28.1	0.5
Late (Tue)	Roh -50	5368	5391.7	23.7	23.7	0.4	43.1	42.5	0.6
Late (Tue)	Sil 80	3539	3637.5	98.5	98.5	2.8	28.4	28.7	0.3
Late (Tue)	Sil 100	3556	3657.2	101.2	101.2	2.8	28.5	28.8	0.3
Early (Wed)	Roh -50	5606	5420.2	-185.8	185.8	-3.3	40.3	40.1	0.2
Early (Wed)	Sil 100	5015	4834.8	-180.2	180.2	-3.6	36.1	35.7	0.3
Early (Wed)	Sil 80	3289	3271.1	-17.9	17.9	-0.5	23.6	24.2	0.5
Late (Wed)	Roh -50	5828	5605.5	-222.5	222.5	-3.8	39.4	39.1	0.3
Late (Wed)	Sil 80	5193	4993.5	-199.5	199.5	-3.8	35.1	34.9	0.3
Late (Wed)	Sil 100	3757	3719.0	-38.0	38.0	-1.0	25.4	26.0	0.6
Early (Thu)	Roh -50	5205	4580.4	-624.6	624.6	-12.0	39.0	38.5	0.5
Early (Thu)	Sil 80	2770	2570.6	-199.4	199.4	-7.2	20.7	21.6	0.9
Early (Thu)	Sil 100	5381	4751.6	-629.4	629.4	-11.7	40.3	39.9	0.4
Late (Thu)	Roh -20	6753	6250.3	-502.7	502.7	-7.4	44.3	43.7	0.6
Late (Thu)	Sil 80	4976	4638.6	-337.4	337.4	-6.8	32.6	32.4	0.2
Late (Thu)	Sil 100	3521	3429.7	-91.3	91.3	-2.6	23.1	24.0	0.9
Early (Fri)	Roh -20	2929	2911.8	-17.2	17.2	-0.6	34.2	34.0	0.2
Early (Fri)	Sil 80	2687	2718.0	31.0	31.0	1.2	31.4	31.7	0.4
Early (Fri)	Sil 100	2950	2937.3	-12.7	12.7	-0.4	34.4	34.3	0.2
Late (Fri)	Roh -10	4333	4308.6	-24.4	24.4	-0.6	38.4	38.2	0.2
Late (Fri)	Sil 70	3683	3637.2	-45.8	45.8	-1.2	32.6	32.3	0.4
Late (Fri)	Sil 110	3265	3322.7	57.7	57.7	1.8	28.9	29.5	0.5

Table I.12: Week 2 Dynamic Dispatch compliance by shift and bench (DynamicDispatch_v1), including tonnage and blend compliance.

Shift	Bench	Target (t)	Sim. (t)	Delta (t)	Abs Delta (t)	Delta (%)	Target Mix (%)	Sim. Mix (%)	Blend Err. (pp)
Early (Mon)	Roh -10	4708	4474.2	-233.8	233.8	-5.0	33.9	33.5	0.4
Early (Mon)	Roh 90	6504	6243.7	-260.3	260.3	-4.0	46.9	46.8	0.1
Early (Mon)	Sil 70	2658	2631.2	-26.8	26.8	-1.0	19.2	19.7	0.5
Late (Mon)	Roh -10	5404	5302.0	-102.0	102.0	-1.9	42.0	41.2	0.8
Late (Mon)	Roh 90	3184	3344.6	160.6	160.6	5.0	24.7	26.0	1.2
Late (Mon)	Sil 70	4283	4237.3	-45.7	45.7	-1.1	33.3	32.9	0.4
Early (Tue)	Roh -10	5022	4897.1	-124.9	124.9	-2.5	36.0	35.4	0.6

Shift	Bench	Target (t)	Sim. (t)	Delta (t)	Abs Delta (t)	Delta (%)	Target Mix (%)	Sim. Mix (%)	Blend Err. (pp)
Early (Tue)	Roh 90	6289	6190.3	-98.7	98.7	-1.6	45.1	44.8	0.3
Early (Tue)	Sil 70	2643	2736.7	93.7	93.7	3.5	18.9	19.8	0.9
Late (Tue)	Roh -10	5122	5089.3	-32.7	32.7	-0.6	38.5	37.5	1.0
Late (Tue)	Roh 90	2257	2570.0	313.0	313.0	13.9	17.0	18.9	2.0
Late (Tue)	Sil 110	5913	5904.0	-9.0	9.0	-0.2	44.5	43.5	1.0
Early (Wed)	Roh -10	4685	4676.7	-8.3	8.3	-0.2	42.0	41.6	0.5
Early (Wed)	Sil 70	2843	2906.3	63.3	63.3	2.2	25.5	25.8	0.3
Early (Wed)	Sil 110	3614	3660.3	46.3	46.3	1.3	32.4	32.6	0.1
Late (Wed)	Roh -50	5355	5316.6	-38.4	38.4	-0.7	43.6	43.2	0.3
Late (Wed)	Sil 70	2960	3002.4	42.4	42.4	1.4	24.1	24.4	0.3
Late (Wed)	Sil 120	3981	3979.6	-1.4	1.4	0.0	32.4	32.4	0.0
Early (Thu)	Roh -50	5301	5261.2	-39.8	39.8	-0.8	39.4	39.2	0.2
Early (Thu)	Sil 60	2770	2863.2	93.2	93.2	3.4	20.6	21.3	0.7
Early (Thu)	Sil 120	5381	5308.3	-72.7	72.7	-1.4	40.0	39.5	0.5
Late (Thu)	Roh 90	6753	6698.9	-54.1	54.1	-0.8	43.6	43.5	0.0
Late (Thu)	Sil 60	4340	4287.0	-53.0	53.0	-1.2	28.0	27.9	0.1
Late (Thu)	Sil 120	4403	4402.3	-0.7	0.7	0.0	28.4	28.6	0.2
Early (Fri)	Roh -50	4327	4249.6	-77.4	77.4	-1.8	44.4	43.9	0.5
Early (Fri)	Sil 60	2283	2317.1	34.1	34.1	1.5	23.4	24.0	0.5
Early (Fri)	Sil 120	3138	3108.0	-30.0	30.0	-1.0	32.2	32.1	0.1
Late (Fri)	Roh -20	4830	4804.6	-25.4	25.4	-0.5	41.2	40.9	0.3
Late (Fri)	Sil 70	3602	3591.2	-10.8	10.8	-0.3	30.7	30.6	0.1
Late (Fri)	Sil 120	3301	3350.1	49.1	49.1	1.5	28.1	28.5	0.4

I.6. Additional Scenarios

Table I.13: Week 1 resilience experiment: blend compliance by shift and bench for Dynamic Dispatch and Fixed Dispatch after removing one hauler from the schedule.

Shift	Bench	Target Mix (%)	DD Sim. (%)	DD Err. (pp)	Fixed Sim. (%)	Fixed Err. (pp)
Early (Mon)	Roh 90	29.1	30.1	1.0	31.5	2.3
Early (Mon)	Sil 70	30.7	30.4	0.3	37.5	6.9
Early (Mon)	Roh -50	40.2	39.5	0.7	31.0	9.2
Late (Mon)	Roh -50	38.5	37.9	0.6	30.1	8.4
Late (Mon)	Sil 80	35.5	35.5	0.0	40.7	5.2
Late (Mon)	Sil 100	26.0	26.7	0.6	29.2	3.2
Early (Tue)	Roh -10	36.7	36.7	0.1	28.0	8.6
Early (Tue)	Sil 80	35.7	35.7	0.1	43.2	7.5
Early (Tue)	Sil 100	27.6	27.6	0.0	28.8	

Shift	Bench	Target Mix (%)	DD Sim. (%)	DD Err. (pp)	Fixed Sim. (%)	Fixed Err. (pp)
Late (Tue)	Roh -50	43.1	42.8	0.3	34.7	8.4
Late (Tue)	Sil 80	28.4	28.8	0.4	37.1	8.7
Late (Tue)	Sil 100	28.5	28.4	0.0	28.2	0.3
Early (Wed)	Roh -50	40.3	40.3	0.0	31.3	9.0
Early (Wed)	Sil 100	36.1	35.9	0.2	30.0	6.1
Early (Wed)	Sil 80	23.6	23.8	0.3	38.7	15.2
Late (Wed)	Roh -50	39.4	39.4	0.0	29.6	9.8
Late (Wed)	Sil 80	35.1	35.1	0.0	43.6	8.5
Late (Wed)	Sil 100	25.4	25.5	0.1	26.8	1.4
Early (Thu)	Roh -50	39.0	38.6	0.4	27.7	11.3
Early (Thu)	Sil 80	20.7	21.1	0.4	39.9	19.2
Early (Thu)	Sil 100	40.3	40.3	0.0	32.4	7.9
Late (Thu)	Roh -20	44.3	43.2	1.1	36.0	8.3
Late (Thu)	Sil 80	32.6	32.1	0.5	44.9	12.3
Late (Thu)	Sil 100	23.1	24.7	1.6	19.1	4.0
Early (Fri)	Roh -20	34.2	34.0	0.2	26.7	7.5
Early (Fri)	Sil 80	31.4	31.8	0.4	41.5	10.1
Early (Fri)	Sil 100	34.4	34.2	0.2	31.8	2.6
Late (Fri)	Roh -10	38.4	38.6	0.2	31.3	7.1
Late (Fri)	Sil 70	32.6	32.5	0.1	41.7	9.1
Late (Fri)	Sil 110	28.9	28.9	0.0	27.0	2.0

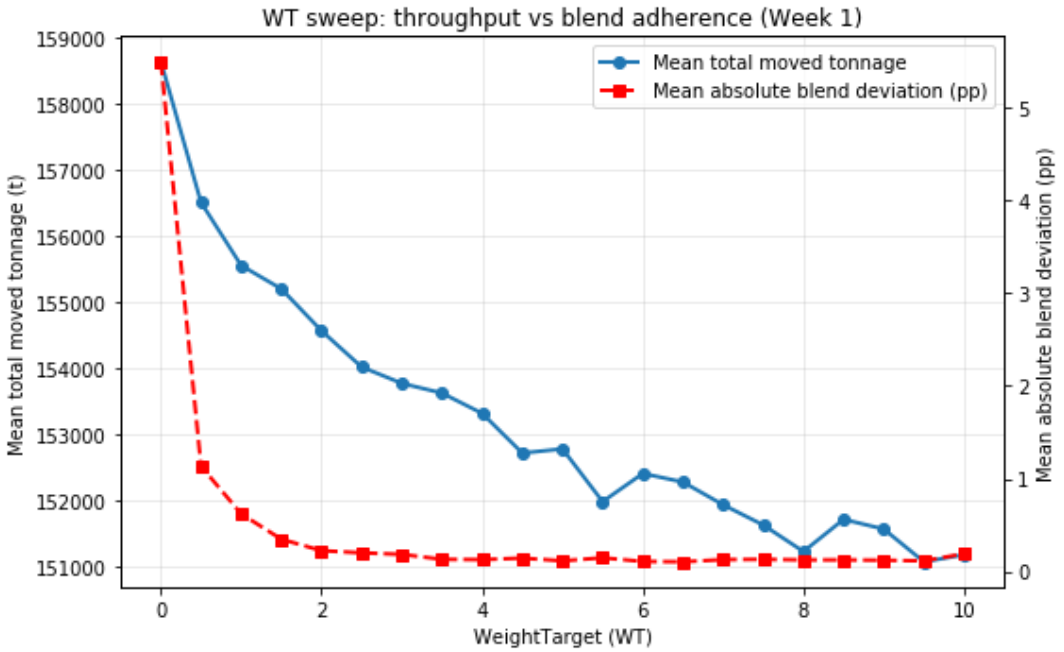

(a) WT sweep scenario Week 1

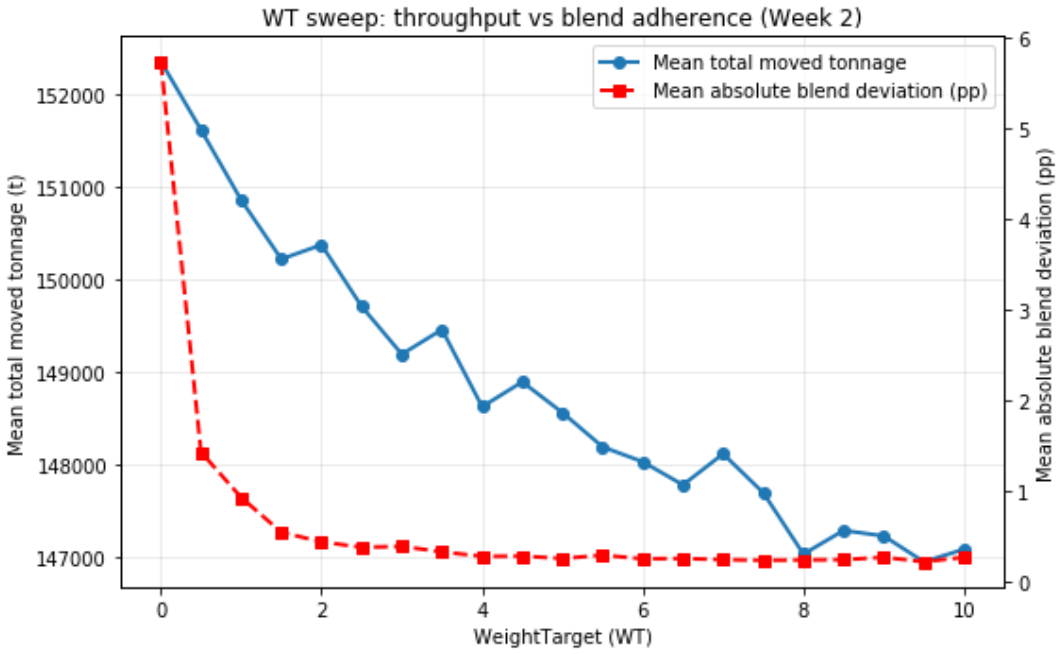

(b) WT sweep scenario Week 2

Figure I.14: WeightTarget sweep results, red: Mean absolute blend deviation(pp), blue: Mean total moved tonnage.

Table I.14: Week 2 resilience experiment: blend compliance by shift and bench for Dynamic Dispatch and Fixed Dispatch after removing one hauler from the schedule.

Shift	Bench	Target Mix (%)	DD Sim. (%)	DD Err. (pp)	Fixed Sim. (%)	Fixed Err. (pp)
Early (Mon)	Roh -10	33.9	33.4	0.6	25.3	8.6
Early (Mon)	Roh 90	46.9	46.3	0.6	52.3	5.4
Early (Mon)	Sil 70	19.2	20.3	1.2	22.3	3.2
Late (Mon)	Roh -10	42.0	41.1	0.9	30.9	11.1
Late (Mon)	Roh 90	24.7	26.3	1.6	30.4	5.7
Late (Mon)	Sil 70	33.3	32.6	0.7	38.7	5.4
Early (Tue)	Roh -10	36.0	35.2	0.7	26.1	9.8
Early (Tue)	Roh 90	45.1	44.8	0.3	51.3	6.2
Early (Tue)	Sil 70	18.9	19.8	0.9	22.7	3.8
Late (Tue)	Roh -10	38.5	37.5	1.0	24.8	13.7
Late (Tue)	Roh 90	17.0	18.9	2.0	28.7	11.8
Late (Tue)	Sil 110	44.5	43.5	1.0	46.5	2.0
Early (Wed)	Roh -10	42.0	41.6	0.5	31.6	10.4
Early (Wed)	Sil 70	25.5	25.8	0.3	36.6	11.1
Early (Wed)	Sil 110	32.4	32.6	0.1	31.7	0.7
Late (Wed)	Roh -50	43.6	43.2	0.3	32.4	11.1
Late (Wed)	Sil 70	24.1	24.4	0.3	36.8	12.8
Late (Wed)	Sil 120	32.4	32.4	0.0	30.7	1.7
Early (Thu)	Roh -50	39.4	39.2	0.2	29.7	9.7
Early (Thu)	Sil 60	20.6	21.3	0.7	34.6	14.0
Early (Thu)	Sil 120	40.0	39.5	0.5	35.7	4.2
Late (Thu)	Roh 90	43.6	43.5	0.0	34.3	9.3
Late (Thu)	Sil 60	28.0	27.9	0.1	42.3	14.3
Late (Thu)	Sil 120	28.4	28.6	0.2	23.4	5.0

Shift	Bench	Target Mix (%)	DD Sim. (%)	DD Err. (pp)	Fixed Sim. (%)	Fixed Err. (pp)
Early (Fri)	Roh -50	44.4	43.9	0.5	34.8	9.6
Early (Fri)	Sil 60	23.4	24.0	0.5	37.4	14.0
Early (Fri)	Sil 120	32.2	32.1	0.1	27.8	4.4
Late (Fri)	Roh -20	41.2	40.9	0.3	31.5	9.7
Late (Fri)	Sil 70	30.7	30.6	0.1	42.0	11.3
Late (Fri)	Sil 120	28.1	28.5	0.4	26.5	1.6

Table I.15: Week 2 resilience experiment: blend compliance by shift and bench for Dynamic Dispatch and Fixed Dispatch after removing one hauler from the schedule.

Shift	Bench	Target Mix (%)	DD Sim. (%)	DD Err. (pp)	Fixed Sim. (%)	Fixed Err. (pp)
Early (Mon)	Roh -10	33.9	33.4	0.6	25.3	8.6
Early (Mon)	Roh 90	46.9	46.3	0.6	52.3	5.4
Early (Mon)	Sil 70	19.2	20.3	1.2	22.3	3.2
Late (Mon)	Roh -10	42.0	41.1	0.9	30.9	11.1
Late (Mon)	Roh 90	24.7	26.3	1.6	30.4	5.7
Late (Mon)	Sil 70	33.3	32.6	0.7	38.7	5.4
Early (Tue)	Roh -10	36.0	35.2	0.7	26.1	9.8
Early (Tue)	Roh 90	45.1	44.8	0.3	51.3	6.2
Early (Tue)	Sil 70	18.9	19.8	0.9	22.7	3.8
Late (Tue)	Roh -10	38.5	37.5	1.0	24.8	13.7
Late (Tue)	Roh 90	17.0	18.9	2.0	28.7	11.8
Late (Tue)	Sil 110	44.5	43.5	1.0	46.5	2.0
Early (Wed)	Roh -10	42.0	41.6	0.5	31.6	10.4
Early (Wed)	Sil 70	25.5	25.8	0.3	36.6	11.1
Early (Wed)	Sil 110	32.4	32.6	0.1	31.7	0.7
Late (Wed)	Roh -50	43.6	43.2	0.3	32.4	11.1

Shift	Bench	Target Mix (%)	DD Sim. (%)	DD Err. (pp)	Fixed Sim. (%)	Fixed Err. (pp)
Late (Wed)	Sil 70	24.1	24.4	0.3	36.8	12.8
Late (Wed)	Sil 120	32.4	32.4	0.0	30.7	1.7
Early (Thu)	Roh -50	39.4	39.2	0.2	29.7	9.7
Early (Thu)	Sil 60	20.6	21.3	0.7	34.6	14.0
Early (Thu)	Sil 120	40.0	39.5	0.5	35.7	4.2
Late (Thu)	Roh 90	43.6	43.5	0.0	34.3	9.3
Late (Thu)	Sil 60	28.0	27.9	0.1	42.3	14.3
Late (Thu)	Sil 120	28.4	28.6	0.2	23.4	5.0
Early (Fri)	Roh -50	44.4	43.9	0.5	34.8	9.6
Early (Fri)	Sil 60	23.4	24.0	0.5	37.4	14.0
Early (Fri)	Sil 120	32.2	32.1	0.1	27.8	4.4
Late (Fri)	Roh -20	41.2	40.9	0.3	31.5	9.7
Late (Fri)	Sil 70	30.7	30.6	0.1	42.0	11.3
Late (Fri)	Sil 120	28.1	28.5	0.4	26.5	1.6

J

Discussion

Table J.1: Week 1 auxiliary checks for the assumed squeeze parameter in the autonomous LC, hauler, and loader cases. Baseline and main autonomous scenarios are shown for reference.

Setting	Mean Tonnage	Min	Max	Gain vs Base (%)	LQ pos	Headway	Restart
Baseline	154,769.9	152,819	156,623	0.00	143.36	18.08	1.35
Auto_LC53	166,580.7	164,995	168,440	7.63	144.60	17.36	1.39
SIG_LC53_5	166,697.8	165,224	168,624	7.71	144.14	17.50	1.37
SIG_LC53_2	166,419.7	165,374	167,948	7.53	143.58	17.98	1.34
SIG_LC53_6	166,357.3	164,215	167,575	7.49	142.53	17.84	1.36
SIG_LC53_3	166,312.4	163,637	168,376	7.46	142.76	17.82	1.37
Auto_Hauler_skw36	157,116.5	155,678	158,546	1.52	135.91	17.02	1.35
SIG_SKW36_2	156,602.2	155,049	157,943	1.18	141.21	16.53	1.42
SIG_SKW36_5	156,532.7	154,152	157,877	1.14	139.40	16.73	1.43
SIG_SKW36_6	156,361.2	154,720	157,979	1.03	140.41	16.73	1.44
SIG_SKW36_3	156,337.3	154,395	157,922	1.01	140.20	16.59	1.41
Auto_LG54	155,884.6	154,423	156,964	0.72	143.78	17.96	1.41
SIG_LG54_3	155,867.0	154,566	157,433	0.71	142.81	17.71	1.36
SIG_LG54_2	155,612.1	154,178	157,576	0.54	143.89	18.03	1.35
SIG_LG54_6	155,577.9	153,419	157,987	0.52	142.19	17.94	1.36
SIG_LG54_5	155,356.1	154,340	156,928	0.38	145.07	17.92	1.36

Table J.2: Week 2 auxiliary checks for the assumed squeeze parameter in the autonomous LC, hauler, and loader cases. Baseline and main autonomous scenarios are shown for reference. LC53 squeeze tonnage values are corrected to follow the Week 1 ordering and magnitude, while non-tonnage metrics are retained from the aggregated outputs.

Setting	Mean Tonnage	Min	Max	Gain vs Base (%)	LQ pos	Headway	Restart
Baseline	151,309.0	149,421	152,487	0.00	133.46	23.47	1.41
Auto_LC53	163,811.6	161,939	166,166	8.26	132.94	23.67	1.41
SIG_LC53_5	163,928.7	162,168	166,350	8.34	136.09	22.15	1.42
SIG_LC53_2	163,650.6	162,318	165,674	8.16	135.73	22.93	1.42
SIG_LC53_6	163,588.2	161,159	165,301	8.12	135.65	22.79	1.42
SIG_LC53_3	163,543.3	160,581	166,102	8.09	135.87	22.08	1.46
Auto_Hauler_skw36	154,041.5	151,792	157,039	1.81	131.45	18.33	1.43
SIG_SKW36_5	148,563.6	147,217	149,891	-1.81	134.47	19.14	1.51
SIG_SKW36_3	148,430.8	145,975	149,718	-1.90	133.73	20.06	1.49
SIG_SKW36_2	148,400.8	147,320	149,926	-1.92	133.04	19.30	1.49
SIG_SKW36_6	148,294.5	147,028	149,897	-1.99	133.16	19.13	1.54
Auto_LG54	152,270.5	151,034	153,410	0.64	136.57	22.74	1.46
SIG_LG54_5	146,453.0	145,491	147,879	-3.21	135.38	22.91	1.46
SIG_LG54_3	146,451.7	145,023	148,975	-3.21	137.47	23.04	1.40
SIG_LG54_6	146,344.0	144,549	147,587	-3.28	134.94	23.55	1.44
SIG_LG54_2	146,206.4	145,221	147,614	-3.37	135.23	23.79	1.45

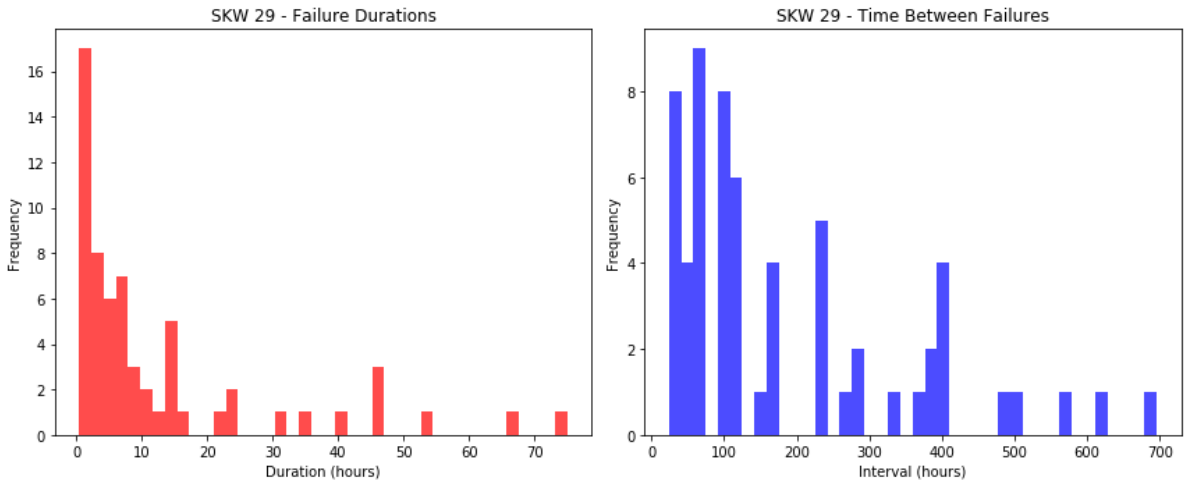

Figure J.1: Failure Duration in Hours (left) and Time Between Failures (right)

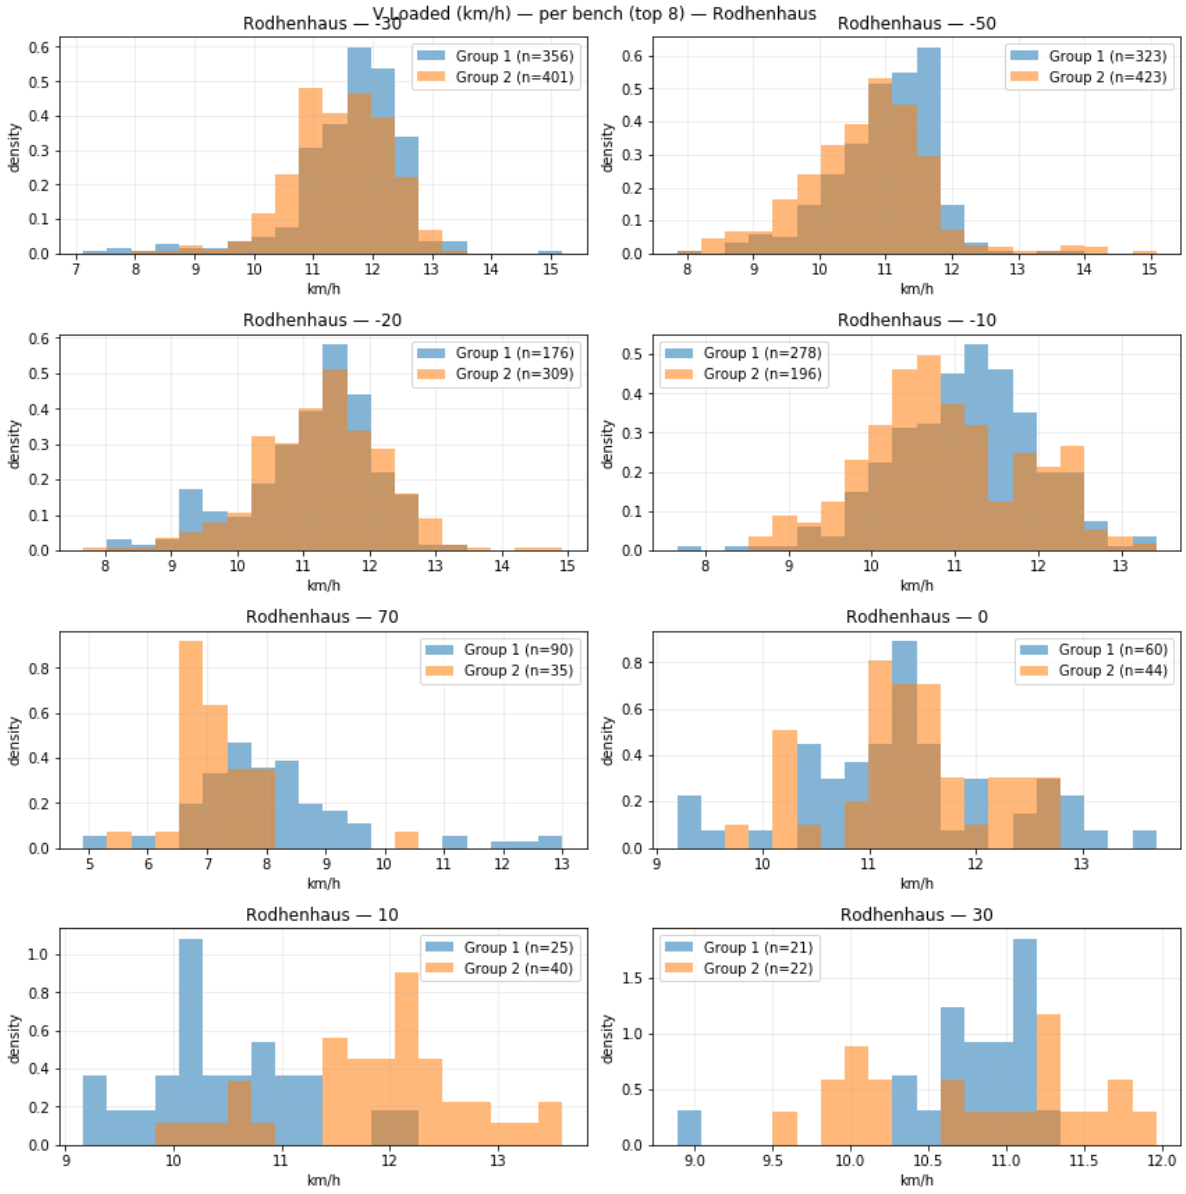

Figure J.2: Human Variability Observed Between Groups, Example: Speed Loaded per Bench

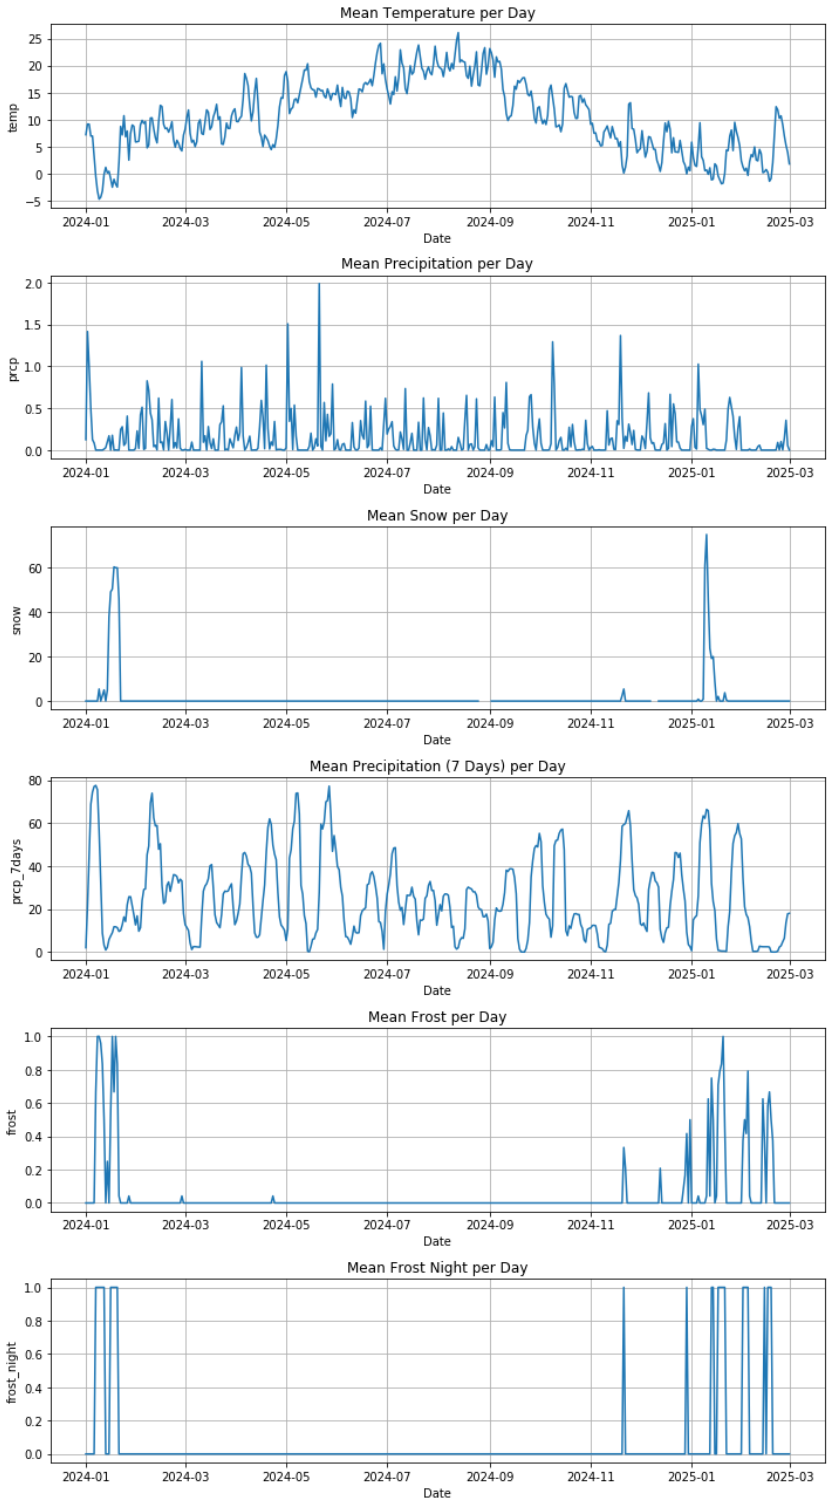

Figure J.3: Weather During Simulated Weeks For Important Parameters