

Katarzyna GDOWSKA
AGH University of Krakow

EFEKTY ZEWNĘTRZNE PROCESU W PROJEKTACH AUTOMATYZACJI: ANALIZA ILUSTRACYJNA NA PRZYKŁADZIE AUTOMATYZACJI PROCE- SU KOMPLETACJI ZAMÓWIEŃ

Streszczenie. Projekty automatyzacji procesów dyskretnych są oceniane względem wewnętrznych celów procesu (czas, koszt, przepustowość), przez co ich skutki dla stron i wartości nieuwzględnionych w funkcji celu – pracowników, osób trzecich czy środowiska – pozostają poza modelem i poza pomiarem. W pracy zastosowano po raz pierwszy do kontekstu automatyzacji zaproponowany niedawno konstrukt *efektów zewnętrznych procesu* (ang. *Process Externalities*). Wkładem jest osadzenie rejestru efektów zewnętrznych w cyklu życia projektu automatyzacji, analityczny wariant kroku ich identyfikacji oraz kompletny, teoretyczny przykład ilustracyjny: automatyzacja kompletacji zamówień (przejście na system *goods-to-person* z algorytmicznym przydziałem zadań). Praca ma charakter koncepcyjny, a walidacja empiryczna stanowi kierunek dalszych badań.

PROCESS EXTERNALITIES IN AUTOMATION PROJECTS: AN ILLUSTRATIVE ANALYSIS OF ORDER-PICKING AUTOMATION

Summary. Automation projects for discrete processes are assessed against the internal objectives of the process (time, cost, throughput), so their effects on parties and values not represented in the objective function – workers, third parties, or the environment – remain outside the model and outside measurement. This paper, for the first time in the automation context, applies the construct of *Process Externalities*, recently proposed in the literature. The contribution is an embedding of the externality register in the lifecycle of an automation project, an analytic variant of the elicitation step, and a complete, theoretical illustrative example: the automation of order picking (the transition to a goods-to-person system with algorithmic task assignment). The paper is conceptual, and empirical validation is a direction for future work.

1. Introduction

The automation of discrete processes is planned, justified, and evaluated against the internal objectives of the process being automated: gains in cycle time, throughput, cost, and quality, monitored with key performance indicators (KPIs) derived from the same objectives [5]. This framing has a structural blind spot: any consequence of automation that falls on a party or a value *not represented* in the objective function lies outside the model and outside what the organisation measures at all. Effects on the wor-

kers who remain in the process, on displaced staff, on the privacy of the people whose behaviour the automated system logs, or on the environment are invisible by construction rather than by accident. Gdowska et. al. [7], introduced a construct that names this class of effects within Business Process Management (BPM): *Process Externalities* (PEs) – consequences of a process for parties or values that its objective function does not represent, generated systematically by optimising the process against that incomplete objective. The construct transposes the economic notion of an externality [11] to the level of an individual process; its closest relative in the automation context is the empirically documented *behavioural externality* of automation, a spillover onto the human worker adjacent to the automated task [2].

While [7] develops the construct for BPM in general, the present paper asks a narrower, engineering-oriented question: *what does it offer to a typical automation project for a discrete process?* We do not extend the construct or report an empirical study; instead we (i) embed its elicitation procedure and externality register in the lifecycle of an automation project, (ii) propose an analytic elicitation variant suited to such projects, and (iii) work through a complete illustrative example – the transition from manual picker-to-parts order picking to a goods-to-person system with autonomous mobile robots (AMRs) and algorithmic task assignment, a canonical object of optimisation whose effects on human work the optimisation objective does not represent [9, 2]. Section 2. recapitulates the construct, Section 3. adapts it to automation projects, Section 4. develops the example, and Section 5. concludes.

2. The Construct of Process Externalities: A Recapitulation

This section summarises the construct after [7], to which the reader is referred for its full development and positioning. A Process Externality is defined there as an unaccounted consequence of a process that affects parties or values not captured by the process objective function; the definition covers technological (real) externalities – effects that change what a third party can do or must bear – and excludes pecuniary effects that merely redistribute value through prices. The construct’s central claim concerns the *generating mechanism*: PEs arise systematically whenever a process is optimised against an objective that encodes only a subset of the values its execution touches, because the omitted values can then be traded away to improve measured performance. The same incompleteness also conceals the effects: monitoring built on the objective’s proxy indicators inherits their blind spots (Goodhart’s law [13]). A PE is thereby distinguished from a process risk, which threatens the process’s *own* goals, and from the behavioural externalities of [2], which land on a participant whose performance the firm already measures and are thus largely internalised.

Two further elements of [7] are used below. First, its two-axis taxonomy classifies each externality by the *locus* of the effect (a process participant, a third party, society, or the environment) and by the *value dimension* at stake: distributive allocation, social relations, individual autonomy, informational privacy, labour and well-being, or environmental load. Second, its normative layer: whether an identified externality is a responsibility concern is assessed through distributive justice [12] – the fairness of how benefits and burdens are shared – and procedural justice [8] – whether the affected have voice and recourse; the assessment feeds one of four *dispositions*: mitigate, promote,

accept with justification, or monitor.

3. Adapting the Construct to Automation Projects

Externalities cannot, in general, be read off the process's own data: an effect on a party for whom the objective defines no event leaves no trace in the event log, so elicitation has to reach beyond what is recorded [7]. The source paper renders elicitation repeatable as a six-step protocol – extract the process and its objective; state the objective's omissions; trace boundary-crossing spillovers; abduce the generating mechanism; anchor each candidate in documented real-world practice; evaluate acceptability through the two justice lenses – and feeds the tracing step primarily with *design fiction* [6], in order to surface effects of processes that do not yet exist.

Our first adaptation concerns this feeder. An automation project differs from the speculative setting in one decisive respect: both poles of the change are specified. The as-is process exists and is observable; the to-be process is documented in the artefacts the project produces – the requirements specification, the simulation model, the dispatching logic. We therefore propose, as the primary feeder for automation projects, an *analytic walkthrough*: for each value dimension of the taxonomy, the analyst compares the as-is and to-be process and asks what the transition adds, removes, or redistributes for each locus class. Design fiction, workshops, and interviews remain complementary feeders for second-order effects of prolonged operation; the anchoring step is retained unchanged.

The protocol's outputs are recorded in an *externality register*, the central artefact of [7]: one entry per candidate, with fields for the locus, value dimension, valence, generating mechanism, lifecycle stage, real-world anchor, the two justice verdicts, a disposition, and an accountable owner. The register deliberately parallels, without merging into, the project's risk register: the two artefacts point in opposite directions, inward at threats to the project's goals and outward at the effects those goals ignore. Our second adaptation maps the externality-aware activities of the source paper onto the phases of an automation project (fig. 1). Three couplings carry most of the weight, corresponding to the points where the objective's incompleteness is introduced, inspected, and either caught or reproduced: at *scoping*, a boundary-crossing prompt enumerates non-participant stakeholders and unpriced values before the business case fixes the objective; at *design*, a specification review interrogates the objective function – in an automation project, a concrete, inspectable artefact – for omissions before the optimisation model is solved; at *monitoring*, the project defines indicators decoupled from the optimisation proxies, since KPIs derived from the objective cannot show what the objective omits. A register maintained pro forma protects no one; its value depends on genuine deliberation with the affected parties, above all the workers of the process being automated.

4. Illustrative Example: Automation of an Order-Picking Process

The example is theoretical and illustrative: it demonstrates that the construct is applicable end-to-end to a typical automation project, not that the listed effects occur in any particular installation; each candidate is anchored in effects documented in the literature. Consider a distribution warehouse operating a manual picker-to-parts process:




Fig. 1. Externality-aware activities in the lifecycle of an automation project, coupled to the externality register (own elaboration; activities and register fields adapted from [7])

pickers traverse the aisles with pick lists, retrieving items for customer orders. The organisation introduces a goods-to-person system: a fleet of AMRs transports mobile racks to stationary picking workstations, and an algorithmic dispatcher assigns pick tasks to stations and sequences robot missions [9]. The project's objective function can be stated, in aggregate form, as

$$\min f = \alpha \cdot T_{\text{pick}} + \beta \cdot C_{\text{op}} \quad \text{s.t.} \quad Q \geq Q_{\min}, \quad (1)$$

where f denotes the value of the objective function, T_{pick} is the expected order-picking time, C_{op} is the operating cost, Q denotes service quality, such as pick accuracy, $Q_{\min}$ is its required minimum level, and α and β are fixed non-negative weights. The exact formulation is immaterial; what matters is the protocol's second step: the objective *encodes* time, cost, and a quality floor, all internal to the process. Absent from (1) are, among others, the working conditions and autonomy of the pickers who remain at the workstations, the informational exposure created by the telemetry the dispatcher requires, the fate of displaced workers, the social fabric of the shop floor, and the energy and material footprint of the robot fleet – each a value the optimisation of (1) *can* trade away without any KPI registering the loss. The analytic walkthrough across the six value dimensions yields the register entries summarised in table 1:

- **E1 (labour–well-being).** In the manual process, pickers control their pace, route micro-decisions, and informal micro-breaks; in the goods-to-person process the dispatcher sets the pace, since a rack waiting at a station is idle capacity that (1) penalises. Work becomes machine-paced; the physical strain of walking is removed

(a *positive* spillover) but partly replaced by repetitive, stationary strain, and any improvement in T_{pick} achievable by compressing worker idle time will be found by the optimiser. Anchor: the transition to automated picking measurably changes human work characteristics [9], and workers' behaviour changes in the mere presence of robotic counterparts [2]. Disposition: *mitigate* (pace limits as design constraints); *promote* for the removed strain.

- **E2 (informational–privacy).** The dispatcher requires per-pick telemetry: which worker, which station, which item, how long. The same data stream doubles, at zero marginal cost, as a worker-surveillance instrument – absent from (1). Once every pick is logged, workers adapt their behaviour to what the log rewards (Goodhart's law [13]). Anchor: the literature on workplace surveillance [1]. Disposition: *mitigate* – purpose limitation, aggregation, and privacy-preserving analysis of the event data [10].
- **E3 (autonomy–procedural).** Task assignment migrates from a human team leader, with whom a worker can argue, to an algorithm, with which they cannot; the to-be process defines no channel of objection, because (1) has no variable for voice. Disposition: *mitigate* – a contestation channel with human review, whose usage becomes a monitored indicator.
- **E4 (relational–social).** Manual picking is cooperative – pickers meet in aisles, exchange informal knowledge, cover for one another – whereas goods-to-person picking isolates each worker at a station facing a robot-fed rack. The social fabric of the shop floor, never a modelled objective, erodes as a by-product of the process logic [2]. Disposition: *monitor*, with rotation and mixed roles as candidate mitigations.
- **E5 (distributive).** The throughput gains accrue to the firm and, through shorter lead times, to customers; the displacement burden falls on low-qualified workers and the local labour market – parties outside the process boundary altogether: those who bear the cost are not those who decide. Under Carroll's layering of obligations [4], lawful redundancies sit above the legal layer but below the espoused ethical norm; the disposition is *accept with justification*, scrutinised procedurally (are affected workers heard? is retraining offered?) rather than presumed.
- **E6 (environmental).** The AMR fleet adds an electricity load, a battery lifecycle with end-of-life waste, and embodied resources; conversely, a partly “dark” warehouse zone needs less lighting and heating – a *positive* externality of the same design decision. Neither term appears in (1) unless energy is priced into C_{op} , and then only at market price, not social cost; this is the one type already routinely quantified through energy and emissions accounting [3]. Disposition: *monitor* (energy accounting); *promote* for the positive term.

4.1. Internalising Externalities in the Automation Model and Project

Identification alone does not internalise a process externality. For the register to affect the automation project, its entries must translate into changes to the optimisation model, the process specification, the system architecture, or the governance of implementation and operation. The appropriate form depends on whether the externality can

Tabela 1

Externality register for the illustrative order-picking automation project (register structure after [7]; entries: own elaboration)

ID	Locus / value dimension	Effect and generating mechanism	Justice verdict	Stage	Disposition
E1	picker; labour–well-being	machine-paced work: station idle time is penalised by the objective, so the optimiser compresses worker-controlled pauses; walking strain removed (+)	distrib. and proced.: fail	design	mitigate (pace limits); promote (+)
E2	picker; informational–privacy	dispatching telemetry doubles as surveillance; workers adapt to what the log rewards	proced.: fail	implementation	mitigate (purpose limitation)
E3	picker; autonomy–procedural	algorithmic task assignment with no channel of objection; no variable for voice	proced.: fail	design	mitigate (contestation channel)
E4	picker, society; relational–social	isolation at stations erodes cooperation and informal knowledge exchange	distrib.: fail	design, implementation	monitor (rotation, mixed roles)
E5	third party, society; distributive	gains accrue to firm and customers; displacement burden on workers and local labour market	distrib.: fail; proced.: conditional	scoping, design	accept with justification (voice, retraining)
E6	environment; environmental	fleet energy load, battery lifecycle (–); less lighting and heating in dark zones (+)	distrib. across generations	operation	monitor (energy account.); promote (+)

be represented by a reliable measure and whether it is directly controlled by the decisions included in the automation model. Consequently, internalisation should not be understood simply as adding all externalities to a single monetary objective function. Let x denote the vector of dispatching, task-assignment, and robot-routing decisions, and let X denote the feasible set defined by the original technical constraints of the goods-to-person system. The baseline problem corresponding to (1) may be written as (2).

$$\min_{x \in X} \alpha T_{\text{pick}}(x) + \beta C_{\text{op}}(x) \quad \text{s.t.} \quad Q(x) \geq Q_{\min} \quad (2)$$

For E1, internalisation changes the feasible set rather than merely assigning a financial penalty to worker strain. Let $\mathcal{K}$ denote the set of pickers, and let $W_k(x)$ represent the workload or work-intensity indicator for picker $k \in \mathcal{K}$, incorporating, for example, pick frequency, uninterrupted workstation time, repetitive task exposure, and the availability of recovery periods. Let $\bar{W}_k$ denote the maximum acceptable workload level established for picker k . The design may then introduce constraints of the form (3).

$$W_k(x) \leq \bar{W}_k, \quad k \in \mathcal{K} \quad (3)$$

These constraints render some throughput-maximising schedules infeasible, preventing the optimiser from improving T_{pick} solely by compressing worker-controlled pauses. Their threshold values should not be selected by the optimisation analyst alone, but es-

established using ergonomic evidence and consultation with the affected workers. At the same time, the reduction in walking and load-carrying effort should be retained as a positive effect and included in the post-implementation assessment rather than treated as cancelling the new stationary and repetitive burden documented for automated order-picking systems [9]. E3 requires a different type of intervention. Procedural autonomy cannot be adequately represented by a weighted penalty in the dispatching objective. It is instead internalised in the to-be workflow by adding an exception path through which a picker can suspend or challenge an assignment and request human review. The specification should define who reviews the decision, the permissible grounds for contestation, and a maximum response time. The resulting process requirement can be monitored through the number of contested assignments, review time, the proportion of decisions modified after review, and repeated patterns of contestation. Internalisation therefore changes not only the parameters of algorithmic assignment, but also the institutional relationship between the worker and the automated decision. For E2, the required intervention is located primarily in the data architecture. The telemetry required for AMR coordination and task dispatching should be separated from data used for employee evaluation. Purpose limitation, restricted access, aggregation, defined retention periods, and privacy-preserving process analysis [10] should be specified before the event-data infrastructure is implemented. The externality register therefore affects which events are recorded, at what level of granularity, for how long they are retained, and which organisational roles can access worker-level data. A technically useful event log need not automatically become a permanent individual performance record. E4 and E5 cannot be internalised fully within the operational dispatching model. E4 may lead to shift-rotation rules, mixed human–technical roles, scheduled team interaction, and the preservation of selected coordination tasks that would otherwise disappear from the to-be process. E5 requires measures at the level of project governance, including early communication, employee participation, transparent workforce planning, retraining, and justification of how the gains and burdens of automation are distributed. These actions do not eliminate displacement, but they affect its distributive and procedural acceptability. In particular, a decision to accept E5 with justification is not equivalent to leaving it unmanaged; it requires the rationale, consultations, mitigating measures, and an accountable decision owner to be documented. Finally, E6 can be partly incorporated into the model by introducing an energy-consumption constraint (4).

$$E_{\text{sys}}(x) \leq \overline{E} \quad (4)$$

In constraint (4), $E_{\text{sys}}(x)$ denotes the net energy consumption of the automated system under decision vector x , and $\overline{E}$ is the maximum acceptable energy-consumption level. Here, $E_{\text{sys}}(x)$ should cover the net system effect, including AMR charging, workstation operation, and any reduction in lighting or heating in automated warehouse zones. Battery replacement, embodied resources, and end-of-life treatment remain lifecycle-level effects and therefore require complementary environmental accounting rather than dispatch-level optimisation alone [3]. The example thus produces four distinct forms of internalisation: constraints imposed on optimisation decisions, redesign of the to-be workflow, restrictions embedded in the information architecture, and organisational measures accompanied by independent monitoring. The indicators derived from the register must remain decoupled from the original performance proxies. Otherwise, a system may

satisfy every target in (2) while E1–E6 remain invisible. The register changes the project only when each accepted disposition is translated into a requirement, an owner, an indicator, and a decision point at which corrective action can still be taken.

5. Discussion, Limitations, and Conclusions

The order-picking example shows that internalising Process Externalities does not require that every social or environmental value be monetised and incorporated into a single enlarged objective function. Some effects can be represented through measurable constraints, as illustrated by workload limits and energy consumption. Others require changes to the workflow, data architecture, or governance of the automation project. This distinction is important because values such as procedural voice, privacy, workplace relations, and the fairness of workforce displacement would be distorted if treated merely as additional operating costs. The analysis also demonstrates that externalities do not arise only from a malfunctioning or improperly implemented system. They may be generated by a system that works exactly as intended. The dispatcher can reduce picking time, improve resource utilisation, and satisfy the required quality level while simultaneously intensifying work, expanding surveillance, reducing opportunities for contestation, or redistributing employment burdens. These effects are systematic because the corresponding values are absent from the criteria used to generate and evaluate the solution. They therefore cannot be detected reliably through conventional project KPIs alone. Moreover, the same automation decision may generate both positive and negative externalities. Goods-to-person picking can reduce walking-related physical strain while increasing repetitive stationary work; it can raise the electricity and material demand of the AMR fleet while reducing lighting and heating requirements in partly automated zones. The externality register prevents these effects from being collapsed into an unsupported overall judgement that automation is either beneficial or harmful. Instead, it retains their distinct loci, mechanisms, and required responses. From a project perspective, the register can function as a boundary artefact connecting process analysts, optimisation specialists, automation engineers, data-protection personnel, human-resource managers, and affected workers. Its main value lies not in producing an exhaustive list of possible concerns, but in linking each identified effect to a specific intervention in the project life-cycle. The example therefore extends the evaluation of automation beyond the question of whether the technical system reaches its internal targets to the question of which unrepresented values make achievement of those targets possible and who bears the resulting burdens.

Three limitations bound the contribution. The analysis is theoretical: although each mechanism is anchored in documented effects [9, 2, 1], we make no claim about the incidence or severity of E1–E6 in any real installation. The protocol yields candidates, not findings, and relies on the analyst’s judgment. And there is no single agreed standard of fairness, so the construct frames the question of whether an externality is acceptable but does not answer it. The immediate agenda is empirical: applying the protocol in a real automation project with inter-analyst agreement as a reliability measure, coupling the register to the simulation and digital-twin models such projects already build, and integrating the decoupled indicators with privacy-preserving process mining [10].

In sum, automation projects for discrete processes are engineered and audited

against internal objectives, so their consequences for unrepresented parties and values are invisible by construction. This paper applied the construct of Process Externalities [7] to this setting for the first time. Its contribution is the adaptation, not the construct: an analytic elicitation variant exploiting the fact that an automation project specifies both the as-is and the to-be process, an embedding of the externality register in the project phases, and an illustrative analysis in which a standard objective was shown to omit six classes of values, each traced to an anchored generating mechanism and assigned a justice verdict and a disposition that changes the project at the design, implementation, or monitoring phase.

The research was financed by a subsidy and the program „Excellence Initiative – Research University” for the AGH University of Krakow.

REFERENCES

1. Ajunwa I., Crawford K., Schultz J.: Limitless worker surveillance. *California Law Review*, 2017, Vol.105, p. 735.
2. Beer R., Qi A., Rios I.: Behavioral externalities of process automation. *Management Science*, 2026, Vol.72(1), p. 575–593.
3. Brocke J., Seidel S., Recker J.: *Green Business Process Management: Towards the Sustainable Enterprise*. Business and Economics. Springer Berlin Heidelberg 2012.
4. Carroll A. B.: The pyramid of corporate social responsibility: Toward the moral management of organizational stakeholders. *Business Horizons*, 1991, Vol.34(4), p. 39–48.
5. Dumas M., Rosa M., Mendling J., Reijers H.: *Fundamentals of Business Process Management*. Springer Berlin Heidelberg 2019.
6. Dunne A., Raby F.: *Speculative Everything: Design, Fiction, and Social Dreaming*. BusinessPro collection. MIT Press 2013.
7. Gdowska K., Suchenia A., Kluza K.: Process externalities: A missing construct for responsible BPM. *BPM 2026 Responsible BPM Process Technology Educators Forum*, 2026,. in press.
8. Leventhal G. S.: What should be done with equity theory? new approaches to the study of fairness in social relationships. *Social exchange: Advances in theory and research* Springer, 1980, p. 27–55.
9. Loske D.: Empirical evidence on human learning and work characteristics in the transition to automated order picking. *Journal of Business Logistics*, 2022, Vol.43(3).
10. Mannhardt F., Koschmider A., Baracaldo N., Weidlich M., Michael J.: Privacy-preserving process mining: F. mannhardt et al. *Business & Information Systems Engineering*, 2019, Vol.61(5), p. 595–614.
11. Pigou A.: *The Economics of Welfare*. Transaction Publishers 2000.

12. Rawls J.: *A Theory of Justice: Revised Edition*. Harvard 2009.
13. Strathern M.: 'improving ratings': audit in the british university system. *European review*, 1997, Vol.5(3), p. 305–321.