*This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026.*
*ISBN 978-952-65771-2-8.*

# Quick Wins From AI Agents in SMEs: Design Science Matrix

## Tobias Schmallenbach *

University of Siegen, Chair of Services Development for SMEs and Crafts, Kohlbettstr 15, 57072 Siegen, Germany
E-mail: t*obias.schmallenbach@uni-siegen.de*

## Matthias Vogel

University of Siegen, Chair of Services Development for SMEs and Crafts, Kohlbettstr 15, 57072 Siegen, Germany.
E-mail *matthias.vogel@uni-siegen.de*

## Prof. Giuseppe Strina

University of Siegen, Chair of Services Development for SMEs and Crafts, Kohlbettstr 15, 57072 Siegen, Germany
E-mail *strina@wiwi.uni-siegen.de*

## Carl Leonard Borgard

University of Siegen, Chair of Services Development for SMEs and Crafts, Kohlbettstr 15, 57072 Siegen, Germany
E-mail *borgardleo@outlook.de*

* Corresponding author

**Abstract:** Small and medium-sized enterprises (SMEs) face strong pressure to adopt artificial intelligence (AI), yet generic adoption guidance rarely tells managers which process to start with or what counts as a near-term win in their specific organisational context. This paper reports a Design Science Research (DSR) study that develops and evaluates an AI-agent quick-win matrix for SMEs. Eight expert interviews across micro, mid-sized and large firms in procurement and sales informed four iterative artefact versions. The final matrix maps five operational processes to feasible agent configurations and rates each by implementation effort and time-to-value. Quotation creation emerges as the most robust cross-context quick-win domain; invoice and tail-spend agents become feasible from the mid-sized segment upward. We theorise a size-contingent logic of agent adoption and offer practitioners a replicable scoping instrument.

**Keywords:** Keywords: AI agents; agentic AI; SMEs; quick wins; design science research; digital innovation; process automation; contingency; expert interviews; quotation automation

*This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.*

## 1 Introduction

Generative AI and the recent class of "AI agents" - systems able to interpret inputs, plan multi-step actions and execute them across digital tools such as e-mail, calendars, document repositories and ERP systems - have shifted the frontier of enterprise automation (Davenport and Ronanki, 2018; Finio and Downie, 2025). Around 30 % of firms worldwide already deploy agent-like systems for business processes (Stanford University and McKinsey, 2024), and the global AI market is forecast to grow from approximately USD 230 billion in 2024 to USD 630 billion by 2028 (IDC, 2024). Surveys among German Mittelstand firms report that 26 % of respondents already use AI agents and that more than half automate inbound e-mail handling with AI (VIER, 2024).

For small and medium-sized enterprises (SMEs), however, this opportunity collides with chronic resource constraints. SMEs typically lack dedicated AI roles, integration architecture and the change-management bandwidth that large firms presuppose (Lindner, 2019; Fischer, 2025). Public discourse adds further confusion: media coverage frequently conflates generative chatbots and autonomous agents (Wolfangel, 2025), and managers are unsure where to start, which processes are suitable, and what organisational prerequisites are required to create value. The result is a managerial decision gap. Without a structured instrument, SMEs risk either underinvesting (missing attainable benefits) or overinvesting (starting with high-complexity initiatives that stall before value materialises).

We address this gap by focusing on quick wins from AI-agent adoption - near-term, tangible benefits achievable with comparatively low implementation effort and a short time-to-value (Kotter, 1995; Kern, 2023). We argue that quick-win potential is contingent on organisational size: in micro firms it depends on lightweight, channel-based solutions and minimal training; in larger organisations value depends on integration, orchestration and governance, which lengthens time-to-value but unlocks scale. The research question is therefore:

> *Which business processes in SMEs show particularly high potential for the use of AI agents, which quick wins can be realised, and how does organisational size influence where such agents can be deployed meaningfully with low implementation effort and short time-to-value?*

We answer this question through a Design Science Research (DSR) study that develops and iteratively evaluates an AI-agent quick-win matrix across eight expert interviews in three size classes and two functional domains. Beyond the artefact itself, we contribute a size-contingent quick-win logic for agentic AI adoption that links process characteristics, organisational infrastructure and feasible agent configurations.

The remainder of the paper is organised as follows. Section 2 frames the study theoretically. Section 3 details the DSR design, the interview protocol and the analytical procedure. Section 4 presents the artefact. Section 5 reports the empirical findings by size class. Section 6 discusses the size-contingent logic, Section 7 derives practical implications, and Section 8 closes with limitations and a future-research agenda.

## 2 Theoretical background

### *2.1 AI agents and agentic systems*

We follow recent industry-academic work in distinguishing AI agents from generative AI tools and from classical robotic process automation (RPA). An AI agent is a software system, typically built on large language models (LLMs), that can perceive inputs, reason over goals, plan multi-step actions and act across digital tools, optionally orchestrating other agents (Microsoft, 2026; Finio and Downie, 2025). Two boundaries matter for our analysis. First, generative chat tools answer prompts but do not act on systems; agents do. Public discourse routinely conflates the two - a recent newspaper test labelled Google Gemini an "agent" although the tested feature was generative (Wolfangel, 2025). Second, RPA executes deterministic scripts on stable user interfaces; agents handle semi-structured inputs and exceptions. We differentiate three modes used throughout the matrix: assistive agents support a human in the loop and surface drafts for approval; autonomous agents execute end-to-end tasks within defined boundaries; orchestrators coordinate several specialised agents toward a shared goal. This mapping resonates with the hybrid-intelligence perspective of Seufert and Meier (2023), who frame value creation in knowledge-intensive settings as a dynamic division of labour between human and AI. Lämmermann et al. (2025) extend this into an "AI-shoring" logic in which agents substitute for scarce skills in labour-constrained SMEs.

Across these modes we observe six use-case primitives - extraction, matching, routing, generation, planning and monitoring - that combine into concrete agent recipes.

### *2.2 SMEs and digital innovation under resource constraints*

The European definition distinguishes micro firms (≤ 9 employees, ≤ EUR 2 m turnover), small firms (≤ 49, ≤ EUR 10 m), medium-sized firms (≤ 499, ≤ EUR 50 m) and large enterprises (Siller and Grausam, 2016). In Germany, 99.3 % of all firms are SMEs and 82.6 % are micro firms (Destatis, 2025), yet together SMEs account for only 40.9 % of gross value added at factor cost - a structural asymmetry that frames any policy discussion of SME productivity (Destatis, 2025). SMEs face well-documented digital-innovation barriers: low IT budgets, fragmented data, limited specialised personnel, informal processes and dependence on a few key people (Ihlau and Duscha, 2019; Peters and Nauroth, 2019; Kick et al., 2024). Digital innovation research has shown that adoption pathways differ qualitatively from those of large enterprises: digital artefacts are generative, editable and re-programmable (Yoo et al., 2010), and digital innovation management is reshaped by the dissolution of clear product-process boundaries (Nambisan et al., 2017). For agentic AI in SMEs, this implies that prescriptions derived from large-enterprise pilots - with their data-engineering teams and AI product owners - rarely transfer without translation. Adoption guidance must therefore be process-level and context-sensitive rather than horizontal.

### *2.3 Quick wins as a managerial concept*

The notion of "quick wins" has roots in change-management theory, where Kotter (1995) emphasises early visible successes to sustain transformation momentum and protect political capital. The construct is widely invoked in implementation science but seldom operationalised. We operationalise quick wins along two dimensions: implementation effort (training and integration depth) and time-to-value (elapsed calendar time from go-live to first measurable benefit). A genuine quick win combines low effort with short time-to-value; conditional quick wins arise where one dimension is favourable but not the other; non-quick-wins combine high effort and long time-to-value. This operationalisation transforms an otherwise vague slogan into a measurable evaluation grid that supports prioritisation across competing AI initiatives. Crucially, we argue that the meaning of a quick win is itself organisationally situated: a one-week chat-bot deployment that frees ten hours per week is a quick win for a two-person workshop, while in a large enterprise an "obvious" quick win frequently includes the secondary benefits of data hygiene and role clarification produced by the implementation work itself.

### *2.4 A contingency perspective on agent adoption*

Information systems and innovation management research has long emphasised that the value of a digital initiative depends on the fit between technology, task and organisational context (Lawrence and Lorsch, 1967; Donaldson, 2001). For agentic AI, three contingencies appear particularly relevant. First, process formalisation - the degree to which a process has codified inputs, steps and outputs - sets the floor below which agent value collapses, because an agent has nothing stable to reason about. Second, data availability - the existence of structured, accessible operational data - bounds what an autonomous agent can decide. Third, integration depth - how many systems an agent must traverse to act - drives implementation effort and therefore time-to-value. Existing maturity models capture analogues of these constructs at firm level (e.g., digital readiness, IT capability) but rarely translate them into process-level prioritisation, and they typically assume large-enterprise contexts. Our matrix targets exactly this gap by pairing process-level diagnosis with size-class differentiation, allowing managers to see which contingency binds their next initiative.

## 3 Research design

### *3.1 Research approach*

We follow the Design Science Research paradigm (Hevner et al., 2004; Peffers et al., 2007; vom Brocke and Maedche, 2019; see also Benner-Wickner et al., 2020, for guidance on DSR in thesis-scale projects), which is well suited to fast-moving technology domains because it produces actionable artefacts evaluated against practical relevance and scientific rigour. Concretely, the project follows the six-step DSRM of Peffers et al. (2007) - (1) problem identification and motivation, (2) definition of solution objectives, (3) design and development, (4) demonstration, (5) evaluation, and (6) communication - while applying the seven guidelines of Hevner et al. (2004) as a quality lens, in particular design as an artefact, problem relevance, design evaluation, research contributions and rigour. We position the artefact as a theory for design and action (Winter and Aier, 2016).

### *3.2 Artefact development*

The artefact is an AI-agent quick-win matrix implemented as a structured comparative table. A first version (V1) was built from a market scan of approximately 30 commercial agent offerings, mapped to relevant SME processes (e.g., Relevance AI, 2025; AskLio, 2026; doubleSlash, 2024). V2 added an explicit agent typology (assistive / autonomous / orchestrator), use-case primitives and integration interfaces. V3 narrowed the scope to the five processes that interview partners could meaningfully evaluate - invoice verification, quotation and invoice creation, customer service and communication, procurement, and demand planning - and added implementation-effort and time-to-value columns. V4, the final artefact, embeds the quick-win evaluation logic, integrates expert-derived agent suggestions and is differentiated by size class.

### *3.3 Empirical evaluation: expert interviews*

Eight experts were interviewed in seven semi-structured sessions between October and November 2025, totalling approximately 4.5 hours of recorded conversation. Sampling followed a purposive logic with three criteria: representation of all three size classes (micro, mid-sized, large), coverage of two complementary functional domains (procurement and sales), and concrete operational responsibility for the relevant processes. Table 1 summarises the sample.

**Table 1** Expert sample

| *ID* | *Size class* | *Industry* | *Role* | *Function* | *Duration (min)* |
|---|---|---|---|---|---|
| E1 | Micro (≤ 9) | Automotive workshop | Owner / managing director | All | 11 |
| E2 | Large (corp.) | Electrical / automotive | Global Commodity Manager Equipment | Procurement | 25 |
| E3 | Large (corp.) | Electrical / automotive | Global Commodity Manager Supplies & Services | Procurement | 47 |
| E4 | Large (corp.) | Electrical / automotive | Sales referent to management | Sales | 26 |
| E5 | Mid-sized | Aluminium processing | Head of inside sales | Sales | 44 |
| E6 | Mid-sized | Aluminium processing | Buyer | Procurement | 48 |
| E7 | Mid-sized | Aluminium processing | Buyer / inside sales | Procurement / Sales | 48 (joint) |
| E8 | Mid-sized | Aluminium processing | Sales project manager | Sales | 33 |

Source: Author's own data.

The interview guide had two parts: an inductive part (status quo, perceived AI potential, repetitive tasks, frustrations, current tools) and a deductive part (process-by-process evaluation against the five matrix processes, with explicit prompts on interfaces, agent mode and quick-win expectations). All interviews were recorded and transcribed using MAXQDA 24, with line-numbered transcripts enabling traceable citation.

### *3.4 Analytical procedure and rigour*

We applied process-tracing logic (Niederberger and Wassermann, 2015): the inductive analysis confirmed problem relevance (relevance cycle), the deductive analysis fed directly into iterative artefact refinement (rigor cycle). Two coding lenses were used simultaneously: a deductive lens guided by matrix dimensions (interfaces, agent mode, effort, time-to-value) and an inductive lens capturing practitioner definitions of quick wins, constraints and risks. To safeguard inferential validity, only agent-process combinations with at least two converging expert assessments entered V4. Where evaluators saw a borderline quick win but explicitly judged the win as realistic in their context, the qualitative judgement overrode the mechanical rating - a deliberate design choice consistent with DSR's utility (rather than truth) criterion (Winter, 2008). Expert consent and anonymisation followed standard ethical practice.

While n = 8 is modest, three considerations support sufficiency for this DSR purpose. First, our aim is artefact validation rather than population inference; ex ante DSR routinely uses small expert samples to demonstrate utility (Hevner et al., 2004). Second, the sample is triangulated across size, function and industry, producing structural rather than statistical variance. Third, consensus on the dominant quick-win domains (quotation creation, order confirmation, invoice verification) emerged across experts who had no prior contact, indicating thematic saturation on the central artefact dimensions.

## 4 The AI-agent quick-win matrix (V4)

The matrix is a three-sheet artefact (one sheet per size class) in which rows describe agent-process combinations and columns capture twelve evaluation dimensions: department, process, example agent (with vendor URL), provider, agent mode (assistive / autonomous / orchestrator), short function description, use-case primitives (extraction, matching, routing, generation, planning, monitoring), required interfaces (e.g., WhatsApp, e-mail, ERP, SAP, Excel fileshare, calendar), implementation effort, time-to-value, derived quick-win rating, quick-win description and current status. Implementation effort is bucketed into low (< 2 days training), medium (< 1 month) and high (> 1 month). Time-to-value is bucketed into immediate (< 3 months), mid-term (< 6 months) and long-term (> 6 months). The quick-win classification follows the deterministic legend in Table 2.

A row therefore reads as a self-contained recommendation a manager can act on. For example, the micro-firm sheet contains a row recommending an assistive WhatsApp bot for quotation creation: the function description specifies that incoming WhatsApp requests are classified, that the bot generates draft offers for defined repair types and that all outgoing offers can optionally be routed for human approval; the interface column lists the WhatsApp business account and a small offer database; effort is rated low, time-to-value immediate, quick-win yes; the status-quo column notes that the owner currently spends roughly 20 % of his working week on customer service and offers, providing a tangible return baseline. Tables 3–5 below distil the three size-class sheets into the operational profile reported back to the firms. The full matrix V4 is held as a supplementary spreadsheet (available from the corresponding author).

**Table 2** Quick-win evaluation logic

| *Implementation effort* | *Time-to-value* | *Quick win* |
|---|---|---|
| Low | Immediate | Yes |
| Low | Mid-term | Conditional |
| Low | Long-term | Conditional |
| Medium | Immediate | Conditional |
| Medium | Mid-term | Conditional |
| Medium | Long-term | No |
| High | Immediate | No |
| High | Mid-term | No |

Source: Author’s own data.

*This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.*

Crucially, expert judgement may override the mechanical rating where the practitioner sees a clear quick win - this design choice acknowledges that effort and time-to-value alone cannot capture context-specific business value (e.g., a one-month integration that immediately removes a daily workload bottleneck).

## 5 Findings

**Table 3** Quick-win profile - micro firm (automotive workshop)

| *Process* | *Example agent* | *Mode* | *Effort* | *Time-to-value* | *Quick win* |
|---|---|---|---|---|---|
| Quotation creation (WhatsApp) | WhatsApp bot - extraction + generation | Assistive | Low | Immediate | Yes |
| Quotation + appointment | WhatsApp bot - extraction + generation + planning | Assistive | Low | Immediate | Yes |
| Appointment booking (telephony) | Telephone bot (e.g., BEK) - extraction + matching | Assistive | Low | Immediate | Yes |
| Invoice verification | - | - | - | - | Not feasible (no formalised process) |
| Demand planning | - | - | - | - | Not feasible (heuristic stock decisions) |

Source: Author's own data.

**Table 4** Quick-win profile - mid-sized firm (aluminium processor)

| *Process* | *Example agent* | *Mode* | *Effort* | *Time-to-value* | *Quick win* |
|---|---|---|---|---|---|
| Invoice verification | Perk - extraction + routing | Autonomous | Low | Immediate | Yes |
| Order confirmation | Zalion - matching | Assistive | Low | Immediate | Yes |
| Quote comparison | AI-First - extraction + matching | Assistive | Low | Immediate | Yes |
| Quotation creation | doubleSlash - extraction + routing | Assistive | Medium | Immediate | Conditional → Yes (expert override) |
| Customer service: delivery dates | Zendesk - extraction | Autonomous | Low | Immediate | Yes |

Source: Author's own data.

**Table 5** Quick-win profile - large enterprise (multinational supplier)

| *Process* | *Example agent* | *Mode* | *Effort* | *Time-to-value* | *Quick win* |
|---|---|---|---|---|---|
| Invoice verification (SAP-integrated) | Perk - extraction + routing | Autonomous | Medium | Immediate | Conditional → Yes (data hygiene) |
| Order confirmation | Zalion - matching | Assistive | Low | Immediate | Yes |

| | | | | | |
|---|---|---|---|---|---|
| Tail-spend negotiation | AskLio - generation | Autonomous | Medium | Mid-term | Conditional → Yes (delegation + clean-up) |
| Quotation creation (SAP-coupled) | doubleSlash - extraction + routing | Assistive | Medium | Immediate | Conditional → Yes (tested) |
| Customer service hotline (24/7) | Orchestrator routing 80 % of calls | Orchestrator | High | Mid-term | Portfolio-level win |

Source: Author's own data.

Across the three sheets, twenty-one agent-process combinations were assessed; thirteen received a yes or expert-override-yes rating, six were rated conditional and two were excluded by the process-formalisation floor. The detailed evidence behind each row follows.

### *5.1 Micro firm: channel-based, plug-and-play wins*

For the automotive workshop (E1; two employees, EUR 130 k turnover), three quick wins emerged, all anchored in channels customers already use rather than back-office systems:

- a WhatsApp bot for quotation creation on standard repair requests (e.g., "oil change, VW Golf 7"), with optional human approval before sending (vendor example: Superchat, 2026);
- a combined WhatsApp bot for quotation and appointment scheduling connected to the workshop calendar; and
- a telephone bot (e.g., BEK Telefon KI, 2026) that books appointments through intelligent calendar matching. This operational framing is consistent with established quotation-process literature that emphasises clarity, standardisation and fast turnaround (Hofstadt, 2019).

All three are rated low effort / immediate time-to-value. The owner reported that customer service and quotation work consumed roughly 20 % of his and 5 % of his employee's weekly time (E1, l. 38–115), providing a clear return baseline. He explicitly endorsed using template-based bot construction kits to test the approach himself rather than commissioning bespoke development (E1, l. 52ff.). Importantly, micro-firm value depends on bypassing rather than penetrating IT systems: if calls and WhatsApp are the channels, the agent must live there. Two further observations are diagnostically important. First, invoice verification has no quick-win potential at this size, because no formalised invoice process exists - confirming the process-formalisation floor introduced in Section 2.4. Second, the owner ruled out sophisticated demand-planning agents because his stock decisions ("does the oil drum still feel full?") are themselves heuristic rather than data-driven (E1, l. 112ff.). Where the underlying decision is heuristic, agentification offers no leverage.

### *5.2 Mid-sized firm: bridging media breaks with assistive agents*

The aluminium processor (E5–E8; ~80 employees, EUR 12–35 m turnover depending on respondent) showed clear potential along four converging lines:

- an invoice-verification agent (e.g., Perk, 2025) that ingests bills, extracts booking data and triggers payment with optional human approval - eliminating the manual paper flow (Zettelflow) and the recurring media breaks (in the sense of Peters and Nauroth, 2019) the team described in detail (E6/E7, l. 95ff.); this resonates with the established invoice-verification literature that flags media breaks and manual data entry as the main error sources in mid-sized firms (Elter, 2018; Zeis and Fiebig, 2023; Siegfarth, 2023);
- an order-confirmation assistant (e.g., Zalion, 2025) that monitors the inbox, matches confirmations against orders and pre-drafts clarifying e-mails to suppliers when prices or quantities deviate. E6 and E7 both flagged this as the single highest-leverage win because it pre-empts downstream invoicing errors and rework (E6, l. 334ff.);

- a quote-comparison agent that ends the current practice of printing offers and marking them by hand (E7, l. 238ff.); and
- an e-mail response agent (e.g., Zendesk, 2025) for delivery-date inquiries, freeing inside sales for complex requests.

The recurring theme is the elimination of media breaks - handovers between paper, e-mail, ERP and Excel that slow throughput and inject errors. The team also described concrete operational consequences of unbroken paper dependency: because delivery notes were not digitised, no one in the relevant department could work from home (E7, l. 145), illustrating how a single paper-based artefact can lock in a much wider organisational rigidity. Notably, two interviewees (E6/E7) explicitly rejected customer-facing bots for inbound delivery-date dialogue, preferring to maintain personal customer relationships (E7, l. 197ff.). This is a methodologically important finding: not every technically feasible automation maps to a perceived quick win, because the customer-relationship value of the human channel forms part of the implicit utility function. The same firm welcomed agentification of internal handovers but rejected it at the external customer interface - a contingency invisible to maturity-model analysis.

A further finding from sales project management (E8) is the latent value of an internal knowledge agent over the firm's intranet that grows with use, allowing engineers to query alloy alternatives, parameters and internal know-how holders. Although E8 rated this as longer-term, he emphasised that even partial deployment would prevent today's frequent loss of orders that are currently rejected for capacity reasons in alloys for which alternatives could be found (E8, l. 84, 102ff.). This case illustrates how an apparently long-term initiative can produce implementation-induced quick wins (here: documented alloy substitutability) before reaching full autonomy.

### *5.3 Large enterprise: orchestration, data hygiene and tail-spend automation*

The corporate experts (E2–E4; multinational electrical and automotive supplier) emphasised three layers of value. First, order confirmation and invoice verification - the same agents proposed for the mid-sized firm - but with conditional quick-win status because of higher integration depth (SAP, ReadSoft, banking interfaces) and compliance constraints. Both nevertheless qualify as quick wins in the experts' judgement: order-confirmation pre-checking shifts error correction from the late invoice stage to the early confirmation stage, where it is "many times less costly" to resolve (E2, l. 34), and the resulting data hygiene compounds over time (E3, l. 72ff.).

Second, a tail-spend agent (e.g., AskLio, 2026) for autonomous negotiation of low-value, high-volume items. E2 reported an active deployment plan for exactly this scope. The agent qualifies as conditional by the mechanical rule (medium effort, mid-term time-to-value) but as a real quick win in practice: the alternative is a structural lose-lose between not negotiating (paying too much) and over-investing buyer time on items below economic relevance - a long-known but rarely solved tail-spend problem (Schaller et al., 2021). The agent both delegates the activity and forces the data clean-up that long-form ERP free-text purchase orders had previously postponed; in E3's words, similar latent demand patterns hide today behind heterogeneous textual descriptions of the same C-parts (E3, l. 90).

Third, quotation generation assisted by SAP-aware agents (e.g., doubleSlash, 2024). E4 had personally tested such an agent and rated it a quick win on the basis of measured time savings and error reduction (E4, l. 38ff.). Notably, the same kind of agent appeared in two roles - sales-side quotation generation and procurement-side supplier-quotation reception (E2, l. 54ff.) - suggesting that agent populations will increasingly negotiate with each other across firm boundaries, with implications for procurement strategy that exceed the scope of this paper.

A consistent pattern emerged across all corporate use cases: implementation often produces implementation-induced quick wins - data clean-up, role clarification and removal of hidden inefficiencies - even when the agent itself takes longer to reach full autonomy. Quick wins, in other words, can be by-products of the readiness work, not only of the running agent. The corporate experts also illustrated the relevance of orchestration: a planned 24/7 customer hotline at one of the corporate sites is designed to route 80 % of approximately 6,000 monthly calls to AI agents and only the residual complex cases to human agents (E3, l. 100), demonstrating how at scale the win shifts from per-task automation to portfolio routing.

### *5.4 The cross-context pattern*

Three patterns hold across size classes:

- Quotation creation is the most robust cross-context quick-win domain. It appears as a quick win in all three size classes, albeit with different agent embodiments (channel bot, assistive ERP-coupled agent, SAP-integrated assistant).
- Invoice and order-related agents are size-dependent. Below the mid-sized threshold, the absence of a structured invoice process forecloses agent value; above it, value scales with volume and integration depth.
- The quick-win definition itself is size-contingent. For E1 it means "low time and money effort"; for E3 it explicitly means "orchestration and system integration that creates transparency". Same word, different operational meaning - a finding with direct theoretical implications (see Section 6).

Demand planning was the only matrix process where no quick win emerged in any size class, because the necessary structured historical data was not available in the studied firms.

## 6 Discussion

Our results support a size-contingent logic of agentic-AI adoption, visualised in Figure 1 and summarised in four propositions (P1–P4) that each flag a distinct value mechanism bound to a size class, with quotation creation as a cross-context anchor.

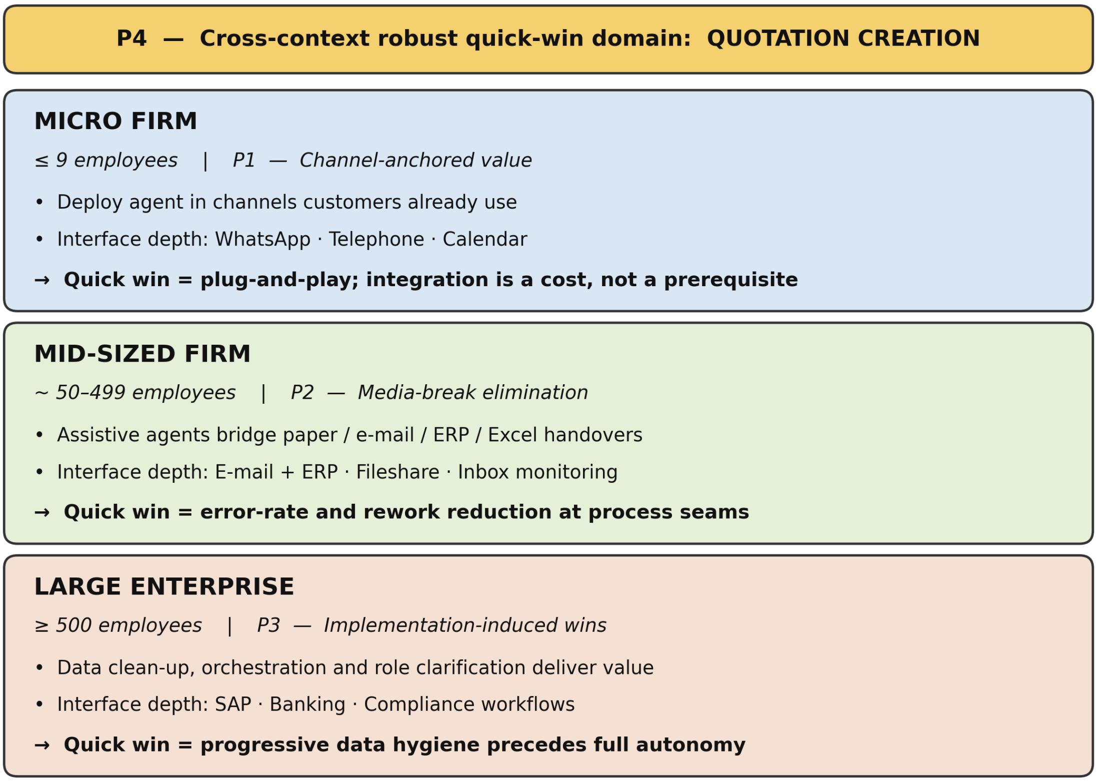


**Figure 1** Size-contingent quick-win logic for AI-agent adoption in SMEs.

- P1 (Channel-anchored value in micro firms). In micro firms, agent value is maximised when agents are deployed in channels customers already use (messaging, telephony) rather than in internal systems. Integration depth is a cost, not a prerequisite.
- P2 (Media-break elimination in mid-sized firms). In mid-sized firms with semi-formalised processes, the largest near-term value comes from assistive agents that bridge media breaks at process handovers (e-mail ↔ ERP, paper ↔ digital), even where the underlying ERP cannot itself be replaced.

- P3 (Implementation-induced quick wins in large firms). In large firms, the readiness work that AI-agent implementation forces - data clean-up, role clarification, tail-spend transparency - frequently delivers quick wins before the agent reaches full autonomy.
- P4 (Cross-context robustness of quotation creation). Among the studied processes, quotation creation is uniquely robust as a quick-win domain across all three size classes, because it combines high-frequency repetition, semi-structured inputs and a natural human-in-the-loop checkpoint at offer release.

This logic refines two strands of prior work. First, against the dominant maturity-model view that treats AI adoption as a linear progression from low to high readiness, our findings show that valuable quick wins exist at every maturity level but along different vectors. Maturity, in other words, conditions which quick win is feasible, not whether one is. Second, the size-contingent meaning of "quick win" itself extends the change-management literature (Kotter, 1995): the definition of a near-term success is endogenous to the organisational context in which it is sought, and a single empirical metric (e.g., hours saved per week) can hide qualitatively different value logics across contexts. This dovetails with calls in digital innovation research to take seriously the editable, generative and re-programmable character of digital artefacts when theorising adoption (Yoo et al., 2010; Nambisan et al., 2017).

The artefact also contributes methodologically. By coupling DSR's iterative artefact refinement with explicit, two-dimensional quick-win operationalisation (effort × time-to-value) and a permissioned override by expert judgement, we provide a replicable ex ante evaluation procedure for fast-moving technology domains where post-implementation measurement is not yet widely available - answering a recent call by vom Brocke and Maedche (2019) for context-sensitive DSR designs. The override mechanism is theoretically meaningful: it acknowledges that effort and time-to-value alone cannot capture business value in domains where the dominant constraint is qualitative (e.g., relationship preservation in the mid-sized customer interface), and it forces the analyst to make that override explicit rather than burying it in the rating.

## 7 Practical implications

For SME owners and functional managers, the matrix supports a staged adoption strategy with three concrete priors. Micro firms should start with channel-native bots in customer-facing processes (WhatsApp, telephony, calendar), retain human approval gates for outgoing offers, and resist deeper integration until customer volume justifies it; the right first question is "where do customers already reach us?", not "which ERP module supports it?". Mid-sized firms should target order confirmation and invoice matching first, framing the proximate win as media-break elimination rather than full autonomy; the order-confirmation agent is particularly attractive because it pre-empts downstream invoicing rework, generating a compounding return. Large firms should deliberately treat the data-readiness work as part of the value case for tail-spend and invoice agents, communicate implementation-induced wins (clean catalogues, surfaced demand patterns, role clarification) to sustain executive sponsorship, and use orchestration to convert per-task gains into portfolio-level routing wins.

For consultants and technology providers, the matrix structures client conversations around three questions: Which channel or system must the agent live in? What effort and time-to-value class does that imply? Which quick-win type applies? This reduces mis-scoping risk and improves expectation management. For investors and analysts, it suggests that vendor evaluation should be size-aware: a vendor whose value proposition assumes deep ERP integration is a poor first match for a micro firm even when the underlying technology is excellent. For policymakers concerned with SME competitiveness, the size-contingent logic suggests that SME-facing AI support programmes should differentiate funding instruments by size and process type - vouchers for plug-and-play deployment in micro firms, integration grants for media-break removal in mid-sized firms - rather than offering a single horizontal subsidy that under-serves both ends of the distribution.

## 8 Limitations and future research

Four limitations bound our claims. First, the sample is purposive and concentrated on procurement and sales in two industries (automotive workshop services and aluminium processing), with one corporate group; HR, R&D, finance and other functions remain unexplored. Second, the evaluation is ex ante and qualitative: we measure perceived utility (Winter, 2008), not realised ROI. Third, the agent market evolves rapidly, so the concrete provider names listed in the matrix are necessarily a snapshot - the contribution rests on the evaluation logic, not on the persistence of any particular vendor. Fourth, sample selection is subject to a self-selection bias: firms that consented to participate showed at least baseline AI-affinity; sceptical firms might have produced systematically different judgements.

Four priorities follow. First, quantitative validation: a follow-on study should track ROI, processing time and error rates pre- versus post-implementation across a larger sample, ideally with matched control units. Second, cross-functional extension: the matrix should be tested in HR and product development, where structured data is sparser but routines exist. Third, theory testing: propositions P1–P4 deserve testing through multi-case comparative designs that vary size and process formalisation as independent variables, ideally with paired pre/post operational metrics. Table 6 operationalises each proposition with a predicted relationship, key moderator and suggested measurement approach, to lower the entry cost for replication teams. Fourth, agent-to-agent dynamics: as both buyer- and supplier-side firms deploy quotation agents, an emerging research question is how negotiation outcomes shift when agents transact with agents - an issue our findings hint at but do not yet resolve. We are currently preparing a longitudinal field study with three implementation partners to begin this programme.

**Table 6** Operationalisation of propositions for empirical replication

| *Proposition* | *Predicted relationship* | *Key moderator* | *Suggested measurement* |
|---|---|---|---|
| P1 | Channel-native deployment → faster time-to-value and higher quick-win realisation in micro firms | Customer channel diversity; integration cost | Days from go-live to first measurable benefit; hours/week saved; pre-post self-report |
| P2 | Assistive media-break-bridging agents → reduced error rate and rework hours in mid-sized firms | Process formalisation; ERP rigidity | Error rate pre/post; rework hours; operator-perceived workload (validated scale) |
| P3 | Implementation work → early ("induced") quick wins preceding full agent autonomy in large firms | Data fragmentation; governance maturity | Data-quality KPIs (duplicate / free-text share); catalogue consolidation; role-clarification survey |
| P4 | Quotation creation remains a quick-win domain across all size classes | Request-repetition rate; input semi-structure | Share of offers bot-drafted; first-pass approval rate; time-to-offer |

Source: Author’s own data.

We close on an ethical note. Several interview partners flagged employee anxiety about job loss as a constraint on agent deployment, consistent with broader survey evidence that 64 % of respondents worldwide and 75 % in Germany fear AI-related job loss (Rittmeier, 2025). A responsible adoption pathway for SMEs must therefore foreground human-in-the-loop design, transparent role redefinition and the redeployment of freed capacity to higher-value work - themes that our matrix, with its assistive-first logic and approval gates, deliberately accommodates.

In sum, the AI-agent quick-win matrix demonstrates that a process-level, size-aware decision instrument can transform a vague managerial promise - "quick wins from AI" - into a structured, evidence-based prioritisation, and that doing so reveals a non-trivial contingency logic of considerable theoretical and practical interest.

## References


AskLio (2026) *Tame the Tailspend - Negotiation Agent*, Available at: https://asklio.ai/ (redirects to https://lio.ai/) (Accessed 10 December 2025).

BEK Telefon KI (2026) *BEK Telefon KI - Intelligente Telefonkommunikation*, Available at: https://bek-agent.de (Accessed 8 December 2025).

Benner-Wickner, M., Kneuper, R. and Schlömer, I. (2020) *Leitfaden für die Nutzung von Design Science Research in Abschlussarbeiten*, IUBH Discussion Papers - IT & Technik, Vol. 2, No. 2, Available at: https://www.econstor.eu/bitstream/10419/229136/1/iubh-dp-it2020-2.pdf (Accessed 23 November 2025).

vom Brocke, J. and Maedche, A. (2019) 'The DSR grid: six core dimensions for effectively planning and communicating design science research projects', *Electronic Markets*, Vol. 29, No. 3, pp. 379–385.

Davenport, T. H. and Ronanki, R. (2018) 'Artificial intelligence for the real world', *Harvard Business Review*, Vol. 96, No. 1, pp. 108–116.

Destatis (2025) *Anteile kleiner und mittlerer Unternehmen an ausgewählten Merkmalen 2023 nach Größenklassen*, Wiesbaden: Statistisches Bundesamt, Available at: https://www.destatis.de/DE/Themen/Branchen-Unternehmen/Unternehmen/Kleine-Unternehmen-Mittlere-Unternehmen/Tabellen/wirtschaftsabschnitte-insgesamt.html (Accessed 28 October 2025).

Donaldson, L. (2001) *The Contingency Theory of Organizations*, Thousand Oaks, CA: Sage.

doubleSlash (2024) *doubleSlash AI Sales Assistant - KI Out of the Box*, Available at: https://www.doubleslash.de/leistungen/ai-assistant/ki-gestuetzte-angebotserstellung/ (Accessed 8 December 2025).

Elter, C. (2018) *Rechnung stellen - Umsatz sichern*, Wiesbaden: Springer Fachmedien. doi:10.1007/978-3-658-18323-3.

Finio, M. and Downie, A. (2025) *Was ist die Orchestrierung von KI-Agenten?*, IBM Think, 26 February, Available at: https://www.ibm.com/de-de/think/topics/ai-agent-orchestration (Accessed 21 November 2025).

Fischer, J. (2025) *Künstliche Intelligenz - Chancen und Herausforderungen für den Mittelstand*, Mittelstand-Digital Zentrum Hamburg, Hochschule für Angewandte Wissenschaften Hamburg, Available at: https://digitalzentrum-hamburg.de/leitfaden/kuenstliche-intelligenz-kmu/ (Accessed 28 December 2025).

Hevner, A. R., March, S. T., Park, J. and Ram, S. (2004) 'Design science in information systems research', *MIS Quarterly*, Vol. 28, No. 1, pp. 75–105.

Hofstadt, P. (2019) *Professionelle Angebotserstellung: Erklärungsbedürftige Inhalte überzeugend verkaufen*, Wiesbaden: Springer Fachmedien. doi:10.1007/978-3-658-27056-8.

IDC (2024) *Künstliche Intelligenz - Umsatz weltweit bis 2028*, Statista, Available at: https://de.statista.com/statistik/daten/studie/1211850/umfrage/umsatz-im-bereich-kuenstliche-intelligenz-weltweit/ (Accessed 25 December 2025).

Ihlau, S. and Duscha, H. (2019) *Besonderheiten bei der Bewertung von KMU: Planungsplausibilisierung, Steuern, Kapitalisierung*, Wiesbaden: Springer Fachmedien. doi:10.1007/978-3-658-18675-3.

Kern, S. (2023) *Das Transformation-Management-Office - Die Basis: Grundlagen und Rüstzeug für agile und digitale Transformationen*, Berlin: Springer. doi:10.1007/978-3-662-68082-7.

Kick, E., Schiedermair, I., Baumgartner, M., Kopp, T. and Kinkel, S. (2024) 'Wissensmanagement in KMU: Identifikation und Priorisierung geschäftskritischer Prozesse', *Zeitschrift für wirtschaftlichen Fabrikbetrieb*, Vol. 119, No. 6, pp. 423–426. doi:10.1515/zwf-2024-1076.

Kotter, J. P. (1995) 'Leading change: Why transformation efforts fail', *Harvard Business Review*, Vol. 73, No. 2, pp. 59–67.

Lämmermann, L., Mayer, V., Schweizer, A. and Urbach, N. (2025) 'AI-Shoring als Zukunftsstrategie: Mit KI-Agenten dem Fachkräftemangel begegnen', *Wirtschaftsinformatik & Management*. doi:10.1365/s35764-025-00564-3.

Lawrence, P. R. and Lorsch, J. W. (1967) *Organization and Environment*, Boston, MA: Harvard Business School Press.

Lindner, D. (2019) *KMU im digitalen Wandel: Ergebnisse empirischer Studien zu Arbeit, Führung und Organisation*, Wiesbaden: Springer Fachmedien. doi:10.1007/978-3-658-24399-9.

Microsoft (2026) *Copilot und KI-Agents*, Microsoft Copilot, Available at: https://www.microsoft.com/de-de/microsoft-copilot/copilot-101/copilot-ai-agents (Accessed 21 November 2025).

Nambisan, S., Lyytinen, K., Majchrzak, A. and Song, M. (2017) 'Digital innovation management: Reinventing innovation management research in a digital world', *MIS Quarterly*, Vol. 41, No. 1, pp. 223–238.

Niederberger, M. and Wassermann, S. (eds) (2015) *Methoden der Experten- und Stakeholdereinbindung in der sozialwissenschaftlichen Forschung*, Wiesbaden: Springer Fachmedien. doi:10.1007/978-3-658-01687-6.

Peffers, K., Tuunanen, T., Rothenberger, M. A. and Chatterjee, S. (2007) 'A design science research methodology for information systems research', *Journal of Management Information Systems*, Vol. 24, No. 3, pp. 45–77. doi:10.2753/MIS0742-1222240302.

Perk (2025) *Rechnungsmanagement*, Available at: https://www.perk.com/de/produkt/rechnungsmanagement/ (Accessed 10 December 2025).

Peters, R. and Nauroth, M. (2019) *Process-Mining: Geschäftsprozesse smart, schnell und einfach*, Wiesbaden: Springer Fachmedien. doi:10.1007/978-3-658-24170-4.

Relevance AI (2025) *Agent Templates - Invoice Processing and Price Quote Automation*, Available at: https://relevanceai.com/agent-templates-tasks/ (Accessed 21 November 2025).

Rittmeier, D. (2025) 'Datengetriebene Technologien definieren unsere Wirtschaft neu', in Hoffmann, C. H. and Hersberger, S. (eds) *Wie die Künstliche Intelligenz die Wirtschaft verändert: Überblick über die theoretischen Grundlagen und Praxisbeispiele entlang von Ökosystemen*, Wiesbaden: Springer Fachmedien, pp. 85–93. doi:10.1007/978-3-658-46839-2_5.

Schaller, M., Fitterer, R., Höller, J., Brandl, R., Kirschner, R. and Preier, S. (2021) *Digital Procurement Survey 2020–2021*, PwC, Available at: https://www.pwc.de/de/strategie-organisation-prozesse-systeme/pwc-digital-procurement-survey-2020-2021.pdf (Accessed 6 December 2025).

Seufert, S. and Meier, C. (2023) 'Hybride Intelligenz: Zusammenarbeit mit KI-Assistenzsystemen in wissensintensiven Bereichen', *HMD Praxis der Wirtschaftsinformatik*, Vol. 60, No. 6, pp. 1194–1209. doi:10.1365/s40702-023-01012-9.

Siegfarth, C. (2023) *Technischer Einkauf: Praxishandbuch für mittelständische Unternehmen*, Freiburg: Haufe Lexware. doi:10.34157/9783648167533.

Siller, H. and Grausam, A. (2016) *Selbstcontrolling für Selbständige und kleine Unternehmen*, Wiesbaden: Springer Fachmedien. doi:10.1007/978-3-658-09363-1.

Stanford University and McKinsey & Company (2024) *Künstliche Intelligenz - Umsatz weltweit bis 2028*, Statista, Available at: https://de.statista.com/statistik/daten/studie/1211850/umfrage/umsatz-im-bereich-kuenstliche-intelligenz-weltweit/ (Accessed 25 December 2025).

Superchat (2026) *Automatisierungen & Chatbots*, Available at: https://www.superchat.de/produkt/automationen (Accessed 8 December 2025).

VIER (2024) *KI-Einsatz in Kundenbetreuung 2024*, Statista, Available at: https://de.statista.com/statistik/daten/studie/1535292/umfrage/einsatz-von-ki-in-kundenbetreuung-mittelstaendischer-unternehmen/ (Accessed 25 December 2025).

Winter, R. (2008) 'Design science research in Europe', *European Journal of Information Systems*, Vol. 17, No. 5, pp. 470–475. doi:10.1057/ejis.2008.44.

Winter, R. and Aier, S. (2016) 'Design science research perspective on artefacts and design principles in enterprise contexts', in Maedche, A., vom Brocke, J. and Hevner, A. (eds) *Designing the Digital Transformation*, Cham: Springer, pp. 469–476.

Wolfangel, E. (2025) 'KI-Agenten: Mr. Bot, wir haben ein Problem', *Die Zeit*, 13 July, Available at: https://www.zeit.de/digital/2025-07/ki-agenten-hype-arbeitsplatz-software-alltag (Accessed 3 January 2026).

Yoo, Y., Henfridsson, O. and Lyytinen, K. (2010) 'The new organizing logic of digital innovation: An agenda for information systems research', *Information Systems Research*, Vol. 21, No. 4, pp. 724–735.

Zalion (2025) *Zalion - AI Procurement Agents*, Available at: https://zalion.ai (Accessed 11 December 2025).

Zeis, A. and Fiebig, H. (2023) *Kommunale Rechnungsprüfung*, Vol. 4, Berlin: Erich Schmidt Verlag. doi:10.37307/b.978-3-503-21217-0.

Zendesk (2025) *Erstklassiger Kundenservice für alle*, Available at: https://www.zendesk.de/lp/nb-support/ (Accessed 10 December 2025).