AI in Restructurings

Efficiency gains from AI in restructurings come not from the tool, but from what has been put into the tool beforehand. Skipping that build-up does not eliminate errors; it merely relocates where they emerge.

by Markulf Behrendt

CHAPTER 1 · STARTING POINT

Three Stages, One Architecture

> *Restructurings are procedurally intensive. That is not a drawback — it is the procedural scrutiny itself that lends them legal stability. Anyone seeking to deploy AI here does not begin with the tool.*

Restructurings are among the most procedurally intensive undertakings in German employment law. They are governed by the participation rights of works council bodies, by formal information and consultation duties, and — at the implementation stage — by a social selection process that is highly prone to error. Defects at any one of these steps can affect the validity of an individual termination, with corresponding consequences for the workforce, the procedure, and the economics of the measure.

This article focuses on the German Mittelstand and on German subsidiaries of international groups linked through tax or corporate-law structures. In this constellation, typical restructurings range from several dozen to several hundred affected jobs, often across one or more sites, regularly involving a local works council, and in a non-trivial number of cases shaped by group-wide directives from abroad. Such directives can collide with German professional and data-protection law — a point we will return to below.

The question of this contribution is at which specific points AI can today contribute to efficiency in such a restructuring without jeopardizing the legal substance of the measure. The answer depends not primarily on the tool, but on how far the build-out of the underlying AI architecture has progressed. The question is therefore not only where AI use can yield efficiency gains, but above all what kind of AI is required for tangible efficiency gains to be possible at all. To anticipate the conclusion: standard AI systems such as ChatGPT, Harvey,

Claude, Copilot, or Gemini can deliver efficiency gains, but a real difference is made only by AI that has been highly specialized for the specific measure. No new AI needs to be "built" for this; it is enough to enrich an existing one with concrete knowledge.

The architecture of such an AI can be described in three stages.

CHAPTER 2 · STAGE ARCHITECTURE

Three Stages of an AI-Supported Restructuring Practice

> *The build-out of an AI-supported advisory practice proceeds in three sequential stages. With each stage, the efficiency potential rises. With each stage, the regulatory requirements rise as well. The order cannot be reversed.*

2.1 Stage 1 — Building the AI Without Client Data

In Stage 1, the AI is prepared without any client data entering the system. The foundation is the world knowledge of a generalist model. On top of that lie three additional layers: legal structural knowledge, specialist legal content from publishers and databases, and — critically — the experiential knowledge of an attorney who has repeatedly led restructurings. Only this fourth layer turns a subsumption engine into an advisory instrument.

This is the central, often overlooked point. An AI that knows only the statute and the commentary literature can formally examine dismissals for operational reasons. But it does not know at which point in a consultation procedure the works council, by experience, will press for clarification. It does not know which personal relationships exist between the participants, or which open or hidden emotions are at play in the negotiations. It does not know when a conciliation committee will calm down and when it will escalate. This knowledge does not arise from reading; it arises from practice.

Matan Grinberg, CEO of the AI tooling company Factory, recently put it concisely: *"the enemy of AI agents is tacit knowledge — things discussed but not written down."* In restructuring practice, this implicit knowledge is particularly dense and particularly consequential. Anyone building Stage 1 is, to a non-trivial extent, working to make this knowledge explicit — to put it in writing, to structure it, to translate it into review templates.

If Stage 1 is built without deep restructuring experience, the AI is nominally ready for use, but of limited substantive value. It then merely reproduces what is already in every textbook. The efficiency gains in Stage 1 arise precisely where the documented experiential knowledge has been made accessible through the tools — in review templates, patterns of argumentation, vetted standard clauses, stored response patterns for typical works council inquiries, and risk-assessment templates.

In Stage 1, no client data is processed. The strict requirements of the GDPR and Section 203 of the German Criminal Code therefore do not apply; under professional law, Stage 1 is governed solely by the duty of care under Section 43 of the Federal Lawyers' Act in the handling of the tool itself.

2.2 Stage 2 — AI With Client Data

In Stage 2, the AI developed in Stage 1 is applied to specific client matters. Contract portfolios, social data, works council agreements, correspondence with the works council — all of this enters the system and is processed there. The efficiency gains in Stage 2 are once again significantly greater than in Stage 1, because it is no longer only the toolchain that improves; the handling of the individual matter improves as such.

Stage 2, however, is subject to a closed catalog of obligations. Required are: processing in EU-hosted models with a complete data processing agreement under Article 28 GDPR; an express written confidentiality undertaking by the AI provider under Section 203(4) sentence 2 of the Criminal Code, going beyond the data processing agreement; a data protection impact assessment; a firm-internal AI policy and documented training of personnel under Article 4 of the EU AI Act; and finally zero data retention at the model level — that is, the contractually secured waiver of any storage of client data beyond the duration of the engagement.

Added to this is what the McKinsey author group around Asin Tavakoli describes as the actual bottleneck of agentic AI: not the models, but the data architecture beneath them. *"Agentic AI scales on strong data,"* writes Tavakoli; eight out of ten companies fail at scaling due to data limitations. Client data must therefore not only enter the system in compliance with the rules, but also be structurally prepared — with clear definitions, clean lineage, and auditable outputs. Only then can an AI work reliably on it. Merely uploading documents into the system is hardly enough for maximum efficiency gains.

There is, in addition, the question of client willingness. Stage 2 presupposes that the client consents to the processing of its data in an AI system. In the Mittelstand and at group subsidiaries, that consent is not a given. Compliance functions of foreign parents sometimes have highly differentiated requirements that go beyond what German law demands; in other cases, they themselves require the use of US cloud services whose conformity with Section 203 of the Criminal Code is not provided for in standard contracts. Both constellations noticeably narrow the practical scope for Stage 2 use by attorneys.

Where the catalog of obligations is satisfied and the client's consent is in place, Stage 2 lifts processing efficiency for a restructuring substantially. In particular, the drafting of individual hearings under Section 102 of the Works Constitution Act, the initial drafting of the mass dismissal notification, the ongoing maintenance of the Whitebook, and responses to works council inquiries during the consultation procedure are accelerated meaningfully.

2.3 Stage 3 — Client Memory Across Engagements

Stage 3 builds on Stage 2. The client data processed in Stage 2 is consolidated in a persistent client memory that could be retained beyond the individual engagement. The consequence would be an advisory profile previously without precedent in restructuring practice: an AI would know what was discussed last year regarding the client's collective bargaining situation, which social plan logic has proven viable for that client, which arguments have already been exchanged, and with what success.

The efficiency potential of Stage 3 is many times greater than that of Stage 2. Engagements would no longer be handled in isolation, but within a continuously growing knowledge base. Personnel changes at the client or at the firm would no longer entail the loss of advisory-relevant contextual knowledge.

The regulatory hurdles for Stage 3 are equally significant. Under data protection law, the burden of justification shifts considerably, because the processing is no longer purpose-bound to the individual engagement but is set up to span purposes. The principle of purpose limitation under the GDPR becomes the central hurdle. To this are added questions of retention and deletion periods, the criminal-law question under Section 203(4) of the Criminal Code, and the professional-law question of matter separation. The technical localization of the AI — at the client or at the attorney — becomes a legal question in its own right.

As matters stand, Stage 3 is a vision. But it shows the direction in which the tension between the technically possible and the legally permissible will move in the coming years.

INFOGRAPHIC · STAGE BALANCE AT A GLANCE

With each stage, efficiency gains rise — and so do the regulatory requirements.

CHAPTER 3 · ANATOMY

Anatomy of a Mittelstand Restructuring

The typical restructuring in the Mittelstand or at a German group subsidiary consists of roughly forty discrete work steps. A smaller share of these is highly strategic — the Whitebook, the negotiation conduct with the works council, the social selection doctrine, the assessment of the conciliation committee, the communication line, the deployment of a transfer company. The larger share is process work: preparing the contractual basis, drafting formally correct notifications, checking consistency across stacks of documents, drafting individual hearing letters.

This split decides where AI can be deployed. Tool-supported acceleration is conceivable wherever recurring, formally describable process work can be modeled. It is not conceivable where strategy, attorney judgment, or personal negotiation conduct form the core.

CHAPTER 4 · PREPARATION

The Employment-Law Baseline Due Diligence

> *Before any major restructuring stands a task that is regularly underestimated in practice: the employment-law baseline due diligence. It is the insurance policy against surprises during implementation.*

The baseline DD reviews the inventory of employment contracts, works council agreements, collective bargaining provisions, bonus and pension commitments, and other collective-law obligations to determine whether obstacles exist that could limit or prevent the planned measure.

An example illustrates the scope. In a 2018 works council agreement on the introduction of an IT system one finds the following sentence: *"Dismissals for operational reasons based on the introduction of this system are excluded."* A clause of this kind — by no means rare in practice — can block a restructuring years later whose economic rationale lies in the automation of activities. If it only becomes visible during the consultation with the works council, the entire measure must be reconceived in substance — at a point at which the client's structures and communications are already aligned to it.

Conventionally, a cursory baseline DD for a Mittelstand constellation requires between 15 and 50 attorney hours. It covers a review of the relevant works council agreements, the relevant group-level works council agreements, the collective bargaining landscape, the typical contract types, and any residual social plan or settlement provisions from earlier restructurings.

With AI support, this preparation can be significantly accelerated to a defensible depth — provided that the architecture built in Stage 1 contains the review templates that an experienced attorney typically works through. At Stage 2, where client documents enter the system, contract portfolios and works council agreements can be searched systematically for risk clauses without the need for a manual review of each individual annex. The attorney's evaluation of the findings — whether a clause is binding in the concrete context, whether it has expired, whether its interpretation is disputed — remains untouched by this.

The baseline DD is thus a concrete application case in which the staged logic of AI deployment becomes visible. Stage 1 enables a faster and more systematic review structure. Stage 2 additionally enables that structure to be applied to the client's actual document inventory. Stage 3 would make it possible to feed the knowledge from each prior review into the next — subject to the regulatory limitations described above.

CHAPTER 5 · APPLICATION

Fields With Robust Efficiency Potential at Stage 2

5.1 Data Extraction From Contract Portfolios

Before the negotiation of a social plan and reconciliation of interests, what is contractually owed must be known. In practice, this typically means a review of several dozen to several hundred employment contracts together with their annexes, bonus policies, company-car arrangements, non-compete clauses, and pension commitments. With AI support, this review can be completed within a time frame that differs markedly from manual processing. Consistency is higher, because the model does not tire.

5.2 Consistency Review Across Document Stacks

The reconciliation of interests, the social plan, the annexes, the mass dismissal notification, the consultation letters, the hearings under Section 102 of the Works Constitution Act, and the accompanying communications add up — even in Mittelstand procedures — to a volume in which consistency can no longer be secured by attention alone. An AI carries out this review with consistent precision and reliably identifies contradictions, semantic shifts, and missed cross-references. Recent case law shows that even limited inconsistencies can raise questions of validity. The efficiency gain here is therefore not only a matter of time, but of quality.

5.3 Initial Draft of the Mass Dismissal Notification

The mass dismissal notification under Section 17 of the Dismissal Protection Act is an example of a document whose formal requirements are sufficiently codified. Section 17(2) and (3) of the Dismissal Protection Act, the Collective Redundancies Directive 98/59/EC, and the relevant case law of the CJEU and the Federal Labor Court — particularly in recent years — provide a self-contained review framework. A properly prompted AI with a stored review template produces, on this basis, a formally complete first draft. Detail questions — calculating the number of employees normally employed, defining the occupational groups, fixing the cut-off dates — are already addressed in the first draft. The need for iteration between the first draft and the signature-ready version diminishes accordingly.

5.4 Individual Hearing Letters Under Section 102 BetrVG

Section 102 of the Works Constitution Act requires the works council to be heard before each individual dismissal. Even in the Mittelstand, these hearings quickly add up to two- or three-digit numbers. Any incomplete or inaccurate hearing renders the individual dismissal invalid under Section 102(1) sentence 3 of the Works Constitution Act. With AI support, mass production from a structured social-data matrix can be generated in significantly less time. The hearings are individually drafted and contain the formally mandatory components. The attorney’s final review then takes place comprehensively, and not merely on a sample basis.

5.5 Consultation Procedure and Whitebook

Out of the works council's notification under Section 17(2) of the Dismissal Protection Act, concrete follow-up questions arise in practice; the prompt and consistent answering of those questions during the reconciliation-of-interests negotiations decisively shapes the course of the procedure. An AI prepared for typical patterns of follow-up questions produces a first draft response to each incoming inquiry. Anyone who handles restructurings regularly knows that compiling answers to works council questions is, in practice, very often the bottleneck of the negotiations and is often surprisingly time-consuming. AI-supported preparation can make a real difference here.

The Whitebook — the strategic preparation document — can be maintained within the same architecture as a continuously updated document. The data underlying the Whitebook can also be drawn upon by the AI when answering works council questions.

Such applications require Stage 2, and thus the full compliance build-out around the AI in use.

CHAPTER 6 · LIMITS

Fields in Which AI Should Not Yet Be Deployed

6.1 Social Selection Doctrine

At the implementation stage, social selection is often the most demanding part of an operationally driven restructuring in substantive terms. It calls not merely for the mechanical application of a points scheme, but for the prior definition of the comparison groups, the determination of the "interchangeability" of activities, the examination of legitimate operational interests, the application of the right to issue instructions, and above all the examination of comparability at the individual level, taking into account specific prior knowledge and training. These judgments are legally complex, internally politically charged, and concern the life circumstances of specific people. They belong to the core work of the attorney.

6.2 Negotiation With the Works Council

Negotiations with the works council — particularly in the conciliation committee — live on presence, pause, tone of voice, and on the understanding that the other side is not representing an abstract position but the interests of the workforce. This dimension cannot be modeled by tools. It belongs to the in-person performance of the attorneys.

6.3 The Attorney's Final Review

Every mass dismissal notification, every reconciliation of interests, every hearing is released or signed by the responsible attorney. That release is the professionally mandated final review under Section 43 of the Federal

Lawyers' Act and the assumption of liability. It is not replaced by AI; it is made easier by it: the review no longer begins with a blank page.

CHAPTER 7 · FEES

Consequence for the Fee Model

Anyone handling a restructuring with the methods described here is faster. Under pure hourly billing, the billable volume thus declines. This implication has so far been discussed less in the market than the tools themselves, even though it is the central question for the fee structure of the coming years.

Steadily increasing hourly rates does not solve the structural problem. The forward-looking answer is a hybrid fee model — for example, flat fees for AI-capable standard work product and hourly rates for the work that remains genuinely attorney-driven, or a fixed fee for the entire engagement, by which 100% of the efficiency-realization risk shifts to the attorney advisor.

This restructuring of fees requires an honest stocktaking of which advisory services are in fact AI-capable and which are not. It also requires the willingness to communicate transparently to clients where efficiency has been generated.

CHAPTER 8 · CONCLUSION

A Balance Sheet

Efficiency gains from AI in restructurings are real. But they are not the product of a tool; they are the product of a prior build-out. Stage 1 — the preparation of the AI without client data — determines the viability of any later use. If Stage 1 is built without deep experiential knowledge, the tools remain substantively empty.

Stage 2 lifts efficiency considerably once again, but is bound to a closed catalog of obligations and to the willingness of the client. In the Mittelstand and at group subsidiaries, that willingness is not a given; foreign group directives can have an additional blocking effect.

Stage 3 — the persistent interweaving of knowledge across engagements — promises the largest efficiency gain and is critical from a data protection perspective. Its practical availability depends on regulatory developments that remain to be seen.