

Adrian König

From Tariff to Award III

Determinants of Court-Awarded Litigation Costs in Civil Appeals to the Swiss Federal Supreme Court

This paper completes the litigation-cost series at the Swiss apex. Its trilingual corpus contains 24'617 judgments and covers 97.8% of civil appeals disposed of by the Federal Supreme Court from 2011 to 2025. In 1'339 merits appeals with an available value in dispute after three adjudicated exclusions, court-awarded litigation costs compress against their federal tariff baselines, with elasticities of about 0.79 for court fees and 0.88 for party compensation. Party compensation is frequently not awarded and lower-instance costs are revised almost exclusively when the appeal succeeds. The analysis is descriptive.

Category of articles: Academic articles

Areas of law: Civil Procedure Law

Citation: Adrian König, From Tariff to Award III, in: Jusletter 28 September 2026

Contents

- 1. Introduction**
- 2. Legal Framework**
 - 2.1. The Appeal in Civil Matters
 - 2.2. The Federal Cost Regime
 - 2.2.1. Court Costs and the Court Fee
 - 2.2.2. Party Compensation
 - 2.2.3. Tariff Structure
 - 2.2.4. Lower-Instance Costs
 - 2.3. The Course of the Appeal Proceedings
- 3. Data**
 - 3.1. Sample Construction
 - 3.2. Coding and Extraction
- 4. Methodology**
 - 4.1. Variables
 - 4.1.1. Dependent Variables and Tariff Baselines
 - 4.1.2. Case Characteristics
 - 4.1.3. Court-Side Characteristics
 - 4.1.4. Cost Allocation
 - 4.1.5. Procedural Complexity Indicators
 - 4.2. Regression Design
- 5. Results**
 - 5.1. The Shape of Court-Awarded Litigation Costs
 - 5.2. Tariff Compression
 - 5.3. Procedural Effort Within the Band
 - 5.4. The Party-Compensation Margin
 - 5.5. Appeal Outcomes and Legal Subjects
 - 5.6. The Second Cost Level
 - 5.7. Robustness
- 6. An Ex-Ante Perspective**
- 7. Discussion**
- 8. Conclusion**
- Appendices**
 - Appendix 1: Coded Variables**
 - Appendix 2: Supplementary Result Tables**

1. Introduction

[1] Litigation costs comprise court costs and party compensation (Art. 95(1) CPC¹; Art. 65 FSAC²; Art. 68 FSAC). Swiss civil courts set both components under statutory tariffs that establish ranges or other scales and regularly leave room for judicial discretion. At federal level, the governing instruments are the Federal Supreme Court Fee Tariff (FSFT³) and the Federal Supreme Court Party Compensation Regulations (FSPCR⁴). Whether an award follows the applicable tariff is therefore an empirical question because the litigation costs set by a court (hereinafter “court-awarded litigation costs”) need not coincide with the amounts stated by the tariff alone.⁵

[2] This paper is the third part of a research series. Part I examined how court-awarded litigation costs are set at first instance through a study of Swiss commercial courts.⁶ Part II carried the inquiry to cantonal high courts acting on appeal.⁷ In both studies, the tariff served as a point of departure.⁸ Court-awarded litigation costs compressed systematically against tariff baselines.⁹ Procedural characteristics supplied the largest observed contribution beyond the tariff to the explanation of court-cost deviations.¹⁰ Even so, observable characteristics left a large share of the variation unexplained.¹¹

[3] With this paper, the series closes at the apex of the Swiss court system. The Federal Supreme Court. For the analysis pursued in the series, this court is particularly suitable because it offers structural advantages over the cantonal settings. A uniform federal cost framework governs the entire corpus (Art. 65 and Art. 68 FSAC, the FSFT and the FSPCR). Parts I and II instead required canton-specific controls to accommodate heterogeneous cantonal regimes.¹² The court also publishes essentially all its final decisions. The corpus therefore comes close to the full population and the selective-publication caveat that constrains the cantonal studies is materially reduced

1 Swiss Civil Procedure Code (CPC) of 19 December 2008 (SR 272).

2 Federal Supreme Court Act (FSAC) of 17 June 2005 (SR 173.110).

3 Tariff for Court Fees in Proceedings before the Federal Supreme Court (Federal Supreme Court Fee Tariff, FSFT) of 31 March 2006 (SR 173.110.210.1).

4 Regulations on Party Compensation and Compensation for Court-Appointed Representation in Proceedings before the Federal Supreme Court (Federal Supreme Court Party Compensation Regulations, FSPCR) of 31 March 2006 (SR 173.110.210.3).

5 See ADRIAN KÖNIG, From Tariff to Award: Determinants of Court-Awarded Litigation Costs in Swiss Commercial Courts. Jusletter (11 May 2026) (cit. KÖNIG, Part I), para. 2 and ADRIAN KÖNIG, From Tariff to Award II: Determinants of Court-Awarded Litigation Costs in Civil Appeals before Swiss Cantonal High Courts. Jusletter (10 August 2026) (cit. König, Part II), para. 1.

6 KÖNIG, Part I (Fn. 5), paras. 3 f.

7 KÖNIG, Part II (Fn. 5), paras. 3 and 5 f.

8 KÖNIG, Part I (Fn. 5), para. 6; KÖNIG, Part II (Fn. 5), paras. 1 and 7.

9 On the compression finding, see KÖNIG, Part I (Fn. 5), para. 150 and KÖNIG, Part II (Fn. 5), para. 67. For party compensation, Part I reads its own estimates as indicative: without the structurally distinct copyright cases its compensation slope loses significance (KÖNIG, Part I [Fn. 5], paras. 131 and 196). The series statement rests primarily on the court-side component, measured as court costs in Parts I and II and as the court fee here.

10 KÖNIG, Part I (Fn. 5), paras. 7 and 123; KÖNIG, Part II (Fn. 5), paras. 7 and 78.

11 KÖNIG, Part I (Fn. 5), para. 9; KÖNIG, Part II (Fn. 5), paras. 78 f.

12 KÖNIG, Part I (Fn. 5), para. 74; KÖNIG, Part II (Fn. 5), paras. 15 f. and 62.

here.¹³ Further, the observed docket spans all three main court languages, German, French and Italian.

[4] Against this background, this paper asks how far the uniform federal cost framework and observable case and procedural characteristics explain court-awarded litigation costs at the Swiss Federal Supreme Court. It studies pecuniary appeals in civil matters¹⁴ decided by the Federal Supreme Court from 1 January 2011 to 31 December 2025. These appeals challenge decisions of the authorities designated in Art. 75 FSCA. As at the cantonal appellate level, court-awarded litigation costs arise at two levels. The first comprises the costs that the Federal Supreme Court sets for its own proceedings. The second comprises lower-instance costs and compensation, which the court confirms, sets aside or modifies depending on the outcome (Art. 67 and Art. 68(5) FSCA).¹⁵ The main analysis concerns the first level.

[5] The empirical basis is a trilingual corpus covering about 98% of the Federal Supreme Court's civil-appeal dispositions. The analysis uses a calibrated, LLM-assisted legal information extraction process to recover the court-awarded amounts and explanatory characteristics from the judgments.¹⁶ For each cost component, it constructs a tariff baseline and uses the deviation of the award from that baseline as the dependent variable in its principal regressions. The paper adopts an explanatory ex-post perspective and, where the data permit, an ex-ante perspective. Throughout, "determinants" denotes descriptive associations.

[6] The results show that court-awarded litigation costs compress at the Federal Supreme Court as well. Among positive awards, the typical amount rises less than proportionally with its tariff anchor, with implied level elasticities of about 0.79 for the court fee and 0.88 for party compensation. Both estimates are significant under a wild-cluster bootstrap by presiding judge. Within the tariff band, the award's position is associated with observable procedural effort, most clearly with whether the court invited a response. Party compensation is frequently not awarded at all. This no-award margin is closely linked to whether the opposing party was invited to respond. At the second cost level, the allocation of lower-instance costs is revised almost exclusively when the appeal succeeds.

13 Even the choice of a suitable object of study at the lower instances is impeded by publication practice. The publication rate of the first-instance civil courts is generally too low to form a viable corpus. Part I therefore focused on the commercial courts. They act as sole cantonal instances, so appeals from their decisions lie directly to the Federal Supreme Court, and their proceedings nevertheless follow the rules of first instance. Publication practice is heterogeneous and incomplete. Zurich yielded the largest published corpus in each predecessor study. In Part I, its estimated coverage was only 61.4%, leaving about 40% of ordinary, non-settled proceedings unpublished. In Part II, no clean publication rate can be formed for it. St. Gallen had to be excluded because of consistently extensive redactions. See KÖNIG, Part I (Fn. 5), paras. 3, 29-33 and 190 and KÖNIG, Part II (Fn. 5), paras. 31-39 and 108 f. See also the author's remarks in the conclusion.

14 *Beschwerde in Zivilsachen*, Art. 72 FSCA.

15 Part II distinguishes the court-awarded litigation costs of the appeal proceedings from the lower-instance costs under review (KÖNIG, Part II (Fn. 5), paras. 14 and 65). The second level here is referred to throughout as the lower-instance costs. Art. 67 FSCA lets the Federal Supreme Court reallocate the costs of the preceding proceedings where it modifies the contested decision.

16 On legal information extraction, see in particular ADRIAN KÖNIG/DANIEL TSCHERAGHSEPEHR/PHILIP PAWLOWSKY/ARNE HOLICKI/MAXIMILIAN L. HAUSER/THIAGO XAVIER/RAPHAEL ZEMP/ARAM A. HESS/LUC WENGER/LI-CHENG CHEN/BENJAMIN BERKOWITSCH/VAKHTANG MARJANIDZE/SHRUTI JHANWAR/NEREA GARCIA MORAN/ENZO NUENER/JELLE BEGEYN/ANĐELIJA MITRAKOVIC/ANTOINE VERMEULEN/SAPANA DANAI/CALEB CATALAN/ROMAN WIXINGER, LEGEX: An Eight-Jurisdiction Benchmark for Legal Review-Table Extraction. In: ICML 2026 Workshop on AI for Law (AI4Law). Seoul, 2026 (cit. KÖNIG et al., LEGEX).

[7] Because the corpus approaches the full population, the findings are not conditioned by an official selection of decisions for publication as they were in the earlier studies.¹⁷ The compression estimates nevertheless remain confined to merits judgments with an available value in dispute and all analyses are descriptive.

[8] The argument proceeds from the legal framework in Section 2 through the data and methodology in Sections 3 and 4 to the results in Section 5. Section 6 then adds the ex-ante prediction. Section 7 discusses the findings and their limitations before Section 8 closes the paper.

2. Legal Framework

[9] This section provides an overview of the legal framework tailored to the empirical analysis.¹⁸ It first explains which disputes and decisions may reach the Federal Supreme Court by an appeal in civil matters, which prior instances precede the federal proceedings and which standard of review the Court applies. It then distinguishes court costs from party compensation and sets out how each cost consequence is assessed and allocated. Finally, it traces the proceedings from the filing of the appeal to the final judgment. Unless indicated otherwise, statutory references in this section are to the FSCA.

2.1. The Appeal in Civil Matters

[10] An appeal in civil matters is subject to several admissibility requirements. Art. 72 FSCA first defines its subject-matter scope. It encompasses decisions in civil matters as well as decisions in debt-enforcement and bankruptcy matters (Art. 72(1) and (2)(a) FSCA). It also encompasses the public-law decisions listed in Art. 72(2)(b) FSCA that are directly connected with civil law. The legal basis of the dispute determines which type of appeal is available.¹⁹

[11] The decisions that may be challenged are governed by Art. 90-93 FSCA. Final decisions conclude the proceedings (Art. 90 FSCA). Partial decisions may be challenged independently if they dispose of claims that can be judged independently or conclude the proceedings for only some of several joined parties (Art. 91 FSCA). Separately notified preliminary and interim decisions on jurisdiction or recusal must be challenged immediately and cannot be challenged later (Art. 92(1) and (2) FSCA). Other separately notified preliminary and interim decisions may be challenged only if either condition in Art. 93(1) FSCA is met. They must be capable of causing irreparable harm (Art. 93(1)(a) FSCA), or allowing the appeal must immediately result in a final decision and thereby avoid significant time or expense for extensive evidentiary proceedings

17 KÖNIG, Part I (Fn. 5), paras. 32 and 190; KÖNIG, Part II (Fn. 5), paras. 31 and 108-109.

18 The FSCA governs the appeal in civil matters. This overview does not purport to be exhaustive. For comprehensive accounts, see in particular Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), *Basler Kommentar. Bundesgerichtsgesetz*, 3rd ed., Basel 2018 (cit. BSK BGG); LAURENT KILLIAS/PAOLA WULLSCHLEGER, § 100 Beschwerde in Zivilsachen ans Bundesgericht, in: Willi Fischer/Fabiana Theus Simoni/Dieter Gessler (eds.), *Kommentierte Musterklagen zur Zwangsvollstreckung und zu den Rechtsmitteln*, vol. V, 2nd ed., Zurich/Basel/Geneva 2022, 219-234, paras. 1-13; MARCO CHEVALIER, *Die Beschwerde in Zivilsachen an das Bundesgericht*, Basel 2009.

19 The procedural route chosen by a cantonal authority is not decisive for this classification. Disputes arising from supplementary insurance to social health insurance illustrate the rule. They are private-law disputes and are therefore, in principle, subject to the appeal in civil matters. BGE 138 I 274, consid. 1.2; BGE 138 III 2, consid. 1.1; KATHRIN KLETT/ELISABETH ESCHER, Art. 72 BGG, in: Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), *Basler Kommentar. Bundesgerichtsgesetz*, 3rd ed., Basel 2018, Art. 72 N 4a f.

(Art. 93(1)(b) FSCA). If a separate appeal is inadmissible or is not lodged, the interim decision may be challenged together with the final decision to the extent that it affects the content of that decision (Art. 93(3) FSCA).

[12] Art. 75 FSCA governs the preceding chain of instances. Decisions of the last cantonal instances may be challenged.²⁰ Decisions of the Federal Administrative Court and the Federal Patent Court are also covered (Art. 75(1) FSCA). A person has standing if he participated in the proceedings before the lower instance or was denied the opportunity to do so, is particularly affected by the contested decision and has a legitimate interest in its annulment or amendment (Art. 76(1) FSCA).

[13] Access to the Federal Supreme Court requires a distinction between pecuniary and non-pecuniary matters.²¹ Art. 74(1) FSCA subjects only pecuniary matters to a minimum value in dispute. Non-pecuniary civil matters are therefore not subject to a value-in-dispute threshold.²² For the pecuniary matters studied here, the minimum value in dispute is CHF 15'000 in employment and tenancy matters and CHF 30'000 in all other matters (Art. 74(1)(a) and (b) FSCA). Below the applicable threshold, Art. 74(2) FSCA permits an appeal only in the exhaustively listed cases. Three exceptions are particularly relevant to the docket studied here. A question of law of fundamental importance may arise (Art. 74(2)(a) FSCA), a federal statute may provide for a sole cantonal instance (Art. 74(2)(b) FSCA), or the contested decision may have been rendered by the Federal Patent Court (Art. 74(2)(e) FSCA).²³ Whether the ordinary threshold is reached is determined by the admissibility value in dispute under Art. 51 FSCA. For appeals against final decisions, the relief sought that remained at issue before the lower instance is decisive (Art. 51(1)(a) FSCA). Art. 51(1)(b) and (c) FSCA specify the applicable bases for partial decisions as well as preliminary and interim decisions. The admissibility value in dispute thus answers only whether the value requirement for access to the Federal Supreme Court is met. It must be distinguished from the value in dispute relevant to costs (Section 2.2).

20 The cantons must generally designate a higher court to decide as an appellate instance. Art. 75(2)(a)-(c) FSCA permits three exceptions: a sole cantonal instance provided for by federal law, a commercial court acting as the sole cantonal instance and a direct action brought before the higher court with the consent of all parties and a value in dispute of at least CHF 100'000. The CPC provides for a sole cantonal instance in specified subject areas (Art. 5 CPC), for commercial courts (Art. 6 CPC), for disputes arising from supplementary insurance to social health insurance (Art. 7 CPC) and for direct actions before the higher court (Art. 8 CPC). On the sole cantonal instance under Art. 7 CPC, see BGer, 4A_58/2025, 1 April 2025, consid. 1.2; BGE 138 III 2, consid. 1.2.2; BGE 138 III 799, consid. 1.1; BSK BGG-KLETT, Art. 75 N 3-6.

21 Non-pecuniary disputes concern rights that, by their nature, cannot be valued in money, do not form part of a person's assets and are not closely connected with a pecuniary legal relationship. Difficulty of valuation alone does not render a matter non-pecuniary. Conversely, a matter is pecuniary if the action ultimately and predominantly pursues an economic objective. It is sufficient that the decision has direct financial consequences or that the economic interest can be assigned a concrete monetary value. A claim for payment is not required. If non-pecuniary and economic elements coexist, the predominant purpose controls. BGE 142 III 145, consid. 6.1; BGE 133 III 490, consid. 3.2; JULIA HÄNNI/LUKAS X. MEYER, Art. 74 BGG, in: Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), Basler Kommentar. Bundesgerichtsgesetz, 3rd ed., Basel 2018, Art. 74 N 4; MICHEL HEINZMANN, Art. 51 BGG, in: Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), Basler Kommentar. Bundesgerichtsgesetz, 3rd ed., Basel 2018, Art. 51 N 9 f.

22 They remain subject to the other admissibility requirements. The subject matter and appealable decision, the prior instance, the appellant's standing, the time limit and the grounds must still satisfy Art. 72, Art. 90-93, Art. 75, Art. 76, Art. 100 and Art. 42 FSCA.

23 On a question of law of fundamental importance where the minimum value in dispute is not reached, see BGE 133 III 645, consid. 2. On an appeal irrespective of value against the decision of a sole cantonal instance under Art. 7 CPC, see BGer, 4A_58/2025, 1 April 2025, consid. 1.2; BGE 138 III 2, consid. 1.2.2; BGE 138 III 799, consid. 1.1; BSK BGG-HÄNNI/MEYER, Art. 74 N 18-25.

[14] The Federal Supreme Court's standard of review is asymmetric. It generally reviews questions of law comprehensively within the statutory grounds of appeal, while reviewing the facts and the assessment of evidence only under narrow conditions.²⁴ The appellant may invoke the violations of law listed in Art. 95 FSCA. The Federal Supreme Court applies the law ex officio (Art. 106(1) FSCA). Once a legal question has been sufficiently introduced into the proceedings, the Court is bound neither by the parties' legal reasoning nor by that of the lower instance.²⁵ The relief sought by the parties nevertheless sets the outer limit of its adjudicative authority (Art. 107(1) FSCA). The appeal must also identify the decisive reasoning of the contested decision and explain concisely how that decision violates the law (Art. 42(2) FSCA). The Court's review therefore does not extend to every conceivable legal question in the manner of a court of first instance.²⁶ A qualified pleading requirement applies to constitutional rights and to cantonal and intercantonal law (Art. 106(2) FSCA). The Federal Supreme Court examines only complaints that are clearly raised and substantiated.²⁷ The position is reversed for facts.²⁸ The Federal Supreme Court generally bases its decision on the facts established by the lower instance (Art. 105(1) FSCA). The appellant must show that a finding is manifestly inaccurate, meaning arbitrary, or is based on a violation of law within the meaning of Art. 95 FSCA and that correcting the defect may affect the outcome of the proceedings (Art. 97(1) FSCA). This binding effect also extends to the result of the assessment of evidence.²⁹ Merely appellatory criticism is insufficient.³⁰ The Federal Supreme Court may correct or supplement the facts ex officio under the conditions of Art. 105(2) FSCA. This power does not relieve the appellant of his duty to state grounds.³¹ If the appeal concerns a decision on provisional measures, legal review is also narrower.³² Under Art. 98 FSCA, the appellant may invoke only the violation of constitutional rights. Such complaints are again subject to the qualified requirements of Art. 106(2) FSCA.

24 In the typical civil chain of instances, review of the facts narrows in particular. An appeal under the CPC permits a party to invoke the incorrect application of the law and the incorrect establishment of the facts (Art. 310(a) and (b) CPC). Under an objection pursuant to the CPC, a complaint concerning the facts is already limited to manifest inaccuracy (Art. 320(a) and (b) CPC). Art. 97 and Art. 105 FSCA then apply before the Federal Supreme Court. The exceptions to the requirement of two cantonal instances under Art. 75(2) FSCA remain applicable. On the Federal Supreme Court's standard of review, see BGE 140 III 86, consid. 2; BGE 135 III 127, consid. 1.5; KILLIAS/WULLSCHLEGER (Fn. 18), § 100 paras. 3-10.

25 BGE 140 III 86, consid. 2; BGE 135 III 397, consid. 1.4; KILLIAS/WULLSCHLEGER (Fn. 18), § 100 paras. 3-6.

26 BGE 140 III 86, consid. 2; BGE 134 III 102, consid. 1.1; KILLIAS/WULLSCHLEGER (Fn. 18), § 100 paras. 9 f.

27 BGE 133 II 249, consid. 1.4.2; BGE 133 III 393, consid. 6; KILLIAS/WULLSCHLEGER (Fn. 18), § 100 paras. 9 f.

28 See the following sentence.

29 BGE 140 III 264, consid. 2.3; BGE 137 III 226, consid. 4.2; KILLIAS/WULLSCHLEGER (Fn. 18), § 100 paras. 3-6.

30 BGE 135 III 127, consid. 1.5; BGE 140 III 264, consid. 2.3; MARKUS SCHOTT, Art. 97 BGG, in: Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), Basler Kommentar. Bundesgerichtsgesetz, 3rd ed., Basel 2018, Art. 97 N 7.

31 BGE 133 IV 286, consid. 6.2; BGE 140 III 86, consid. 2; JOHANNA DORMANN, Art. 105 BGG, in: Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), Basler Kommentar. Bundesgerichtsgesetz, 3rd ed., Basel 2018, Art. 105 N 65 f.

32 See the following two sentences.

2.2. The Federal Cost Regime

[15] The operative part of the judgment determines two legally distinct cost consequences. Court costs, in the sense of the FSCA, are the procedural costs owed to the Swiss Confederation under Art. 65 and Art. 66 FSCA. Party compensation concerns the reimbursement between the parties of the necessary costs of litigation under Art. 68 FSCA. Procedural success is generally decisive for both consequences (Art. 66(1) and Art. 68(1) FSCA). Their statutory bases differ and the amounts are set separately. Court costs are generally charged to the unsuccessful party (Art. 66(1) FSCA). The Federal Supreme Court may allocate them differently or refrain from charging them if the circumstances so warrant (Art. 66(1) FSCA). Unnecessary costs are borne by the person who caused them (Art. 66(3) FSCA). As to party compensation, the Federal Supreme Court determines in its judgment whether and to what extent the unsuccessful party must reimburse the successful party for his necessary costs (Art. 68(1) and (2) FSCA).³³


2.2.1. Court Costs and the Court Fee

Court fee (item 1 FSFT)		Party compensation (Art. 4 FSPCR)	
Value in dispute (CHF)	Fee range (CHF)	Value in dispute (CHF)	Fee range (CHF)
up to 10'000	200-5'000	up to 20'000	600-4'000
over 10'000 to 20'000	500-5'000	over 20'000 to 50'000	1'500-6'000
over 20'000 to 50'000	1'000-5'000	over 50'000 to 100'000	3'000-10'000
over 50'000 to 100'000	1'500-5'000	over 100'000 to 500'000	5'000-15'000
over 100'000 to 200'000	2'000-8'000	over 500'000 to 1'000'000	7'000-22'000
over 200'000 to 500'000	3'000-12'000	over 1'000'000 to 2'000'000	8'000-30'000
over 500'000 to 1'000'000	5'000-20'000	over 2'000'000 to 5'000'000	12'000-50'000
over 1'000'000 to 5'000'000	7'000-40'000	over 5'000'000	20'000 to 1% ^a
over 5'000'000 to 10'000'000	10'000-60'000		
over 10'000'000	20'000-100'000		

^a Over CHF 5m the party-compensation upper bound is 1% of the value in dispute (lower bound CHF 20'000). The schedules repeat each boundary amount in the two adjacent bands. For reproducibility, the implementation assigns an observation exactly on a boundary to the lower band, consistently with the “up to” presentation above. *Source*: item 1 FSFT (SR 173.110.210.1); Art. 4 FSPCR (SR 173.110.210.3).

Table 1: Federal tariff schedules by value in dispute.

33 Different special rules apply to public entities. The Swiss Confederation, cantons, communes and organisations entrusted with public-law tasks are generally not charged court costs if, in their official sphere and without a pecuniary interest of their own, they apply to the Federal Supreme Court or a decision concerning that sphere is challenged (Art. 66(4) FSCA). If such an entity prevails in its official sphere, it is generally not awarded party compensation (Art. 68(3) FSCA). Unlike Art. 66(4) FSCA, Art. 68(3) FSCA does not additionally require the absence of a pecuniary interest of the entity's own. THOMAS GEISER, Art. 66 und 68 BGG, in: Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), *Basler Kommentar. Bundesgerichtsgesetz*, 3rd ed., Basel 2018, Art. 66 N 26 ff. and Art. 68 N 19-21.



Source: FSFT (SR 173.110.210.1); FSPCR (SR 173.110.210.3); own calculations.

Figure 1: The federal tariff baselines by value in dispute.

[16] Court costs and the court fee are not synonymous. Under Art. 65(1) FSCA, court costs comprise the court fee, the fee for copying written submissions, the specified translation expenses and compensation paid to experts and witnesses. The court fee is therefore only one component of court costs. It is the federal cost component studied here. Its assessment is governed by the FSFT.³⁴ Unlike the cantonal cost regimes examined in Parts I and II, which differ from one another, the FSFT applies to the entire Federal Supreme Court corpus studied here.³⁵

[17] Art. 65(2) FSCA specifies five assessment criteria: the value in dispute, the scope and difficulty of the matter, the manner in which the proceedings are conducted and the parties' financial situation. The value in dispute is otherwise calculated under Art. 51-53 FSCA. Item 1 FSFT gives effect to the value criterion through fee bands, but expressly designates those bands as guidelines. The value in dispute is therefore one statutory criterion among several and does not eliminate the Federal Supreme Court's discretion in assessing the fee.

[18] For an appeal against a final decision, the court fee is based on the economic scope of what remains at issue before the Federal Supreme Court. The relief sought and maintained in the federal proceedings is therefore decisive.³⁶ Parts of the subject matter before the lower instance that are not challenged or are no longer contested are excluded from this assessment. The study terms this basis of assessment the fee value in dispute. It must be distinguished from the admissibility value in dispute. The latter determines under Art. 51 FSCA whether the minimum threshold in Art. 74(1) FSCA is reached and, for final decisions, refers to the relief sought that remained at issue before the lower instance (Art. 51(1)(a) FSCA). A third concept concerns party compensation. Its compensation value in dispute is determined by the relief sought that remains at issue

³⁴ Item 1 FSFT; BSK BGG-GEISER, Art. 65 N 4 f. and 10 f.

³⁵ In Part I, median court costs in the two regression cantons differ by roughly a factor of three and a half, at CHF 15'600 in Zurich compared with CHF 4'290 in Aargau (KÖNIG, Part I [Fn. 5], para. 100). Part II examines five cantonal regimes with different assessment methods (KÖNIG, Part II [Fn. 5], paras. 15 f.). See item 1 FSFT.

³⁶ BGer, 4A_39/2011, 8 August 2011, consid. 14.2; BGer, 2C_874/2017, 12 December 2018, consid. 7.2; BSK BGG-GEISER, Art. 65 N 11.

before the Federal Supreme Court (Art. 3(2) FSPCR).³⁷ The fee value in dispute and the compensation value in dispute thus generally refer to the same procedural stage, but rest on different cost instruments.

[19] For appeals against preliminary and interim decisions, the value in dispute relevant to court costs is determined by the relief sought before the instance hearing the main matter (Art. 51(1)(c) FSCA).³⁸ The scope and difficulty, by contrast, are assessed by reference to the Federal Supreme Court proceedings themselves. In matters involving a pecuniary interest, the court fee generally ranges from CHF 200 to 100'000 under Art. 65(3)(b) FSCA. Within this outer range, the value-based guidelines of the FSFT apply (Table 1). For special reasons, the statutory maximum may be increased to no more than twice its amount (Art. 65(5) FSCA). For the special disputes covered by Art. 65(4) FSCA, the court fee ranges from CHF 200 to 1'000 irrespective of the value in dispute. These include employment disputes with a value in dispute of up to CHF 30'000. Special reasons may increase the upper limit to CHF 10'000 in these cases (Art. 65(5) FSCA).

2.2.2. Party Compensation

[20] Party compensation is not part of court costs. In principle, it is intended to reimburse the successful party for all necessary costs caused by the litigation (Art. 68(1) and (2) FSCA). These include the costs of legal representation and any other necessary costs (Art. 1 FSPCR). The costs of legal representation comprise the lawyer's fee and necessary expenses (Art. 2(1) FSPCR). In matters involving a pecuniary interest, the fee is generally based on the value in dispute (Art. 3(1) FSPCR). Within the prescribed ranges, the Federal Supreme Court considers the importance and difficulty of the matter as well as the scope and time required for the lawyer's work (Art. 3 and Art. 4 FSPCR). It may exceed the tariff range if the workload is exceptional (Art. 8(1) FSPCR). It may reduce the fee if there is a manifest disproportion or if the proceedings end without a decision on the merits (Art. 8 FSPCR). The Federal Supreme Court determines the compensation as a lump sum on the basis of the case file (Art. 12(1) FSPCR). Any value-added tax due is included (Art. 12(1) FSPCR). In practice, however, this tariff-based lump sum is frequently insufficient to cover the successful party's actual legal fees.³⁹

2.2.3. Tariff Structure

[21] The FSFT and the FSPCR are both degressive over the observed range of values in dispute. In their respective component-specific estimation samples, their elasticities with respect to the value in dispute are 0.48 and 0.46, respectively (Section 5.2). The court fee also approaches a statutory ceiling (Figure 1). This tariff structure forms the normative benchmark for the empirical deviation analysis. The Federal Supreme Court nevertheless retains discretion in assessing the amounts within the statutory and regulatory ranges.

37 BSK BGG-GEISER, Art. 68 N 26.

38 BGE 133 III 645, consid. 2.3; BGer, 5A_664/2010, 1 June 2011, consid. 7; BSK BGG-GEISER, Art. 65 N 11.

39 Art. 2(3) and Art. 12(1) FSPCR; BGE 139 III 190, consid. 4.2; BGer, 5A_113/2025, 3 October 2025, consid. 3.5.1; BSK BGG-GEISER, Art. 68 N 4 and 8-11. The FSPCR does not govern the fee owed in the lawyer-client relationship. Its tariff contains no hourly rate and the Court fixes compensation as a lump sum from the file. The cited commentary describes the compensation as frequently insufficient.

2.2.4. Lower-Instance Costs

[22] The costs of the Federal Supreme Court proceedings must be distinguished from the costs of the preceding proceedings. If the Federal Supreme Court modifies the contested decision, it may reallocate the court costs of the preceding proceedings (Art. 67 FSCA). Depending on the outcome of the federal proceedings, it may affirm, set aside or modify the lower instance's decision on party compensation (Art. 68(5) FSCA). It may determine the compensation itself under the applicable federal or cantonal tariff or entrust that determination to the lower instance (Art. 68(5) FSCA). These powers do not alter the fact that the court costs and party compensation set in the Federal Supreme Court's operative part are distinct cost consequences of the appeal proceedings (Art. 65, Art. 66 and Art. 68 FSCA).

2.3. The Course of the Appeal Proceedings

[23] The appeal proceedings begin with the timely filing of the appeal (Art. 100 FSCA). The appeal period is generally 30 days from notification of the complete copy of the contested decision (Art. 100(1) FSCA). The shorter special periods in Art. 100(2)-(4) FSCA concern specific subject areas. The written appeal must contain the relief sought, the grounds with an indication of the evidence and the signature (Art. 42(1) FSCA). The grounds must explain concisely how the contested decision violates the law (Art. 42(2) FSCA). New facts and evidence are admissible only to the extent that the decision of the lower instance gives rise to them. New claims are inadmissible (Art. 99 FSCA). The FSCA does not provide for a cross-appeal.⁴⁰ A party seeking a modification of the contested decision in his favour must therefore lodge an appeal independently and unconditionally within the appeal period.⁴¹

[24] Upon receipt, the Federal Supreme Court examines its jurisdiction *ex officio* (Art. 29(1) FSCA). It also examines the other admissibility requirements *ex officio* and with full power of review.⁴² The judge in charge of the proceedings directs the case until judgment (Art. 32(1) FSCA). As a rule, the appellant must pay an advance equal to the expected court costs (Art. 62(1) FSCA). If the appellant fails to pay within the additional period, the Federal Supreme Court does not entertain the appeal (Art. 62(1) and (3) FSCA). A party without sufficient means may apply for legal aid if his case does not appear devoid of any prospect of success (Art. 64(1) FSCA). Under the simplified procedure, the division president or a judge designated by him disposes of manifestly inadmissible, manifestly insufficiently reasoned, vexatious or abusive appeals in a decision with brief reasons (Art. 108 FSCA). Art. 109 FSCA provides in addition for a summarily reasoned decision by a panel of three judges. It covers, in particular, manifestly unfounded or manifestly well-founded appeals and certain decisions not to entertain an appeal because no question of law of fundamental importance arises (Art. 109(1)-(3) FSCA). If the appeal is withdrawn, the parties settle or the case becomes moot, the judge in charge strikes it from the register (Art. 32(2) FSCA).

⁴⁰ See the following.

⁴¹ BGE 134 III 332, consid. 2.5; BGE 138 V 106, consid. 2.1; BSK BGG-DORMANN, Art. 107 N 2; Federal Council Dispatch on the Total Revision of the Federal Judicature of 28 February 2001, FedGaz. 2001 4202, 4342.

⁴² BGE 133 III 645, consid. 2; BGE 132 III 291, consid. 1; MARKUS BOOG, Art. 29 BGG, in: Marcel A. Niggli/Peter Uebersax/Hans Wiprächtiger/Lorenz Kneubühler (eds.), Basler Kommentar. Bundesgerichtsgesetz, 3rd ed., Basel 2018, Art. 29 N 5 f.

[25] At the same time, interim protection may be required (Art. 103 and Art. 104 FSCA). As a rule, an appeal has no suspensive effect (Art. 103(1) FSCA). In civil matters, however, it has suspensive effect by operation of law within the scope of the relief sought if it is directed against a constitutive judgment (Art. 103(2)(a) FSCA). The judge in charge may modify the suspensive effect ex officio or on application (Art. 103(3) FSCA). He may also order provisional measures to preserve the existing situation or provisionally protect threatened interests (Art. 104 FSCA).

[26] Whether an adversarial exchange of written submissions takes place depends on the individual case (Art. 102(1) FSCA). To the extent necessary, the Federal Supreme Court serves the appeal on the lower instance and on the other parties, participants or authorities entitled to appeal and sets a deadline for a response (Art. 102(1) FSCA). The lower instance submits the case file within that period (Art. 102(1) and (2) FSCA). A response is therefore not required in every case (Art. 102(1) FSCA). As a rule, there is no further formal exchange of written submissions (Art. 102(3) FSCA). This must be distinguished from the right to reply that follows from the right to be heard.⁴³ Each party must receive the submissions of the other participants in the proceedings and have an opportunity to comment on them.⁴⁴ A formal second exchange of written submissions is not required in every case (Art. 102(3) FSCA).

[27] The case is then prepared for adjudication on the basis of the written submissions and the lower instance's case file (Art. 102(1) and (2) and Art. 105(1) FSCA). Art. 55 FSCA permits the Federal Supreme Court to conduct its own evidentiary proceedings. The judge in charge may take the necessary evidence himself or entrust this task to a competent federal or cantonal authority (Art. 55(1) and (2) FSCA). In law, this power remains narrow because the Federal Supreme Court is bound by the facts established by the lower instance and new facts and evidence are restricted (Art. 97(1), Art. 99(1) and Art. 105(1) FSCA). Separately, the division president may order an oral party hearing (Art. 57 FSCA).

[28] Once the matter is ready for adjudication, the competent division generally sits as a panel of three judges (Art. 20(1) FSCA). It generally sits as a panel of five judges for questions of law of fundamental importance or at the request of a judge (Art. 20(2) FSCA). Appeals against decisions of cantonal supervisory authorities in debt-enforcement and bankruptcy matters are excluded from this rule (Art. 20(2) FSCA). Oral deliberations take place if the division president orders them, a judge requests them or unanimity is not reached (Art. 58(1) FSCA). In all other cases, the Federal Supreme Court decides by circulation (Art. 58 FSCA). Subject to the statutory grounds for exclusion, party hearings, oral deliberations and the votes that follow are public (Art. 59(1) and (2) FSCA). The party hearing under Art. 57 FSCA and the panel's oral deliberations under Art. 58 FSCA must therefore be distinguished in conceptual and functional terms (Art. 57, Art. 58 and Art. 59 FSCA).

[29] At the end of the proceedings, the Federal Supreme Court may not go beyond the relief sought by the parties (Art. 107(1) FSCA). If it allows the appeal, it decides the matter itself or remits it to the lower instance for a new decision (Art. 107(2) FSCA). It may also remit the matter to the authority that decided as the first instance (Art. 107(2) FSCA). The operative part deter-

⁴³ See the following.

⁴⁴ Art. 29(2) of the Federal Constitution of the Swiss Confederation (FC) of 18 April 1999 (SR 101); Art. 6(1) of the Convention for the Protection of Human Rights and Fundamental Freedoms (ECHR) of 4 November 1950 (SR 0.101); BGE 133 I 100, consid. 4.3; BGE 137 I 195, consid. 2; BGE 138 I 484, consid. 2.1; BSK BGG-DORMANN, Art. 102 N 24-26.

mines the court costs and, where applicable, party compensation for the Federal Supreme Court proceedings (Art. 66 and Art. 68 FSCA). To the extent required by the outcome, it also reallocates the costs of the preceding proceedings (Art. 67 and Art. 68(5) FSCA). The full written judgment is served on the parties, the lower instance and any other participants (Art. 60(1) FSCA). The judgment becomes final on the day it is rendered (Art. 61 FSCA).

3. Data

[30] This section describes the data. It identifies the corpus and its source, measures the corpus against the court's own case statistics, sets out how the analysis scope and cost-deviation sample are constructed, and reports the coding and LLM-assisted extraction of the variables.

[31] The empirical basis comprises civil judgments from the Federal Supreme Court's two civil-law divisions.⁴⁵ It covers judgments from 1 January 2011 to 31 December 2025. The judgments come from the swiss-caselaw database compiled by the OpenCaseLaw project. The database aggregates officially published decisions into a single, deduplicated, and structured corpus.⁴⁶ The published judgment is the unit of observation.⁴⁷ The study covers all three official languages.⁴⁸

45 The First Civil Law Division and the Second Civil Law Division. The corpus delimitation uses their respective docket prefixes, 4A and 5A (Section 3.1).

46 OpenCaseLaw, *swiss-caselaw* dataset, <https://huggingface.co/datasets/voilaj/swiss-caselaw> (corpus snapshot of 9 June 2026). The dataset republishes officially published judicial decisions, which are not protected by copyright under Art. 5(1)(c) of the Federal Act on Copyright and Related Rights (Copyright Act, CopA) of 9 October 1992 (SR 231.1). Only the records under `court=bger` are used. The parallel `bge` records (the official reports) are a second representation of the same cases and are excluded to avoid double counting.


47 Where the court joins several appeals, it disposes of them in one published judgment, and that judgment is one observation. Thirty-five judgments of the analysis scope are of this kind. In eighteen of them the court imposed a single joint fee, so the one observation carries the effort of several proceedings while its tariff baseline is built from one value in dispute. In the other seventeen it set a fee for each joined appeal separately, and the mismatch does not arise. Four of the thirty-five have since left the cost-deviation sample as purely procedural, so the exclusions below act on seventeen and thirty-one, respectively. Removing the judgments with a joint fee leaves both elasticities unchanged to three decimals, at 0.787 and 0.878. Removing all joined judgments moves them to 0.785 and 0.876 (Appendix 2, R18). The aggregation therefore does not carry the result.

48 The official languages of the Confederation are German, French, and Italian, while Romansh is an official language only in dealings with persons of that language (Art. 70(1) FC). Before the Federal Supreme Court, Rumantsch Grischun is in principle an admissible language of the proceedings (Art. 54(1) FSCA). No Romansh-language judgment appears in the civil-appeal corpus.

Year	Official dispositions	Corpus (4A+5A)	Coverage	Usable	Frame
2011	1'618	1'572	97.2%	1'112	647
2012	1'709	1'675	98.0%	1'210	678
2013	1'631	1'589	97.4%	1'120	602
2014	1'664	1'628	97.8%	1'154	582
2015	1'721	1'685	97.9%	1'169	637
2016	1'685	1'643	97.5%	1'188	658
2017	1'748	1'717	98.2%	1'244	654
2018	1'705	1'662	97.5%	1'209	615
2019	1'654	1'616	97.7%	1'171	633
2020	1'749	1'714	98.0%	1'276	683
2021	1'608	1'576	98.0%	1'175	629
2022	1'709	1'676	98.1%	1'284	652
2023	1'601	1'576	98.4%	1'141	584
2024	1'605	1'579	98.4%	1'182	598
2025	1'754	1'709	97.4%	1'194	574
2011-2025	25'161	24'617	97.8%	17'829	9'426

Coverage: corpus count divided by official dispositions. *Source:* Federal Supreme Court Geschäftsberichte; swiss-caselaw database; own compilation.

Table 2: Corpus coverage against the population of disposed civil appeals.



The cost-deviation sample is the subset of pecuniary merits judgments with an available value in dispute that enters the cost-deviation analysis (Section 3). *Source:* Federal Supreme Court Geschäftsberichte; swiss-caselaw database; own compilation.

Figure 2: Corpus coverage and the cost-deviation sample, by year.

[32] The published judgments studied in Parts I and II may not represent the underlying caseload.⁴⁹ By contrast, the Federal Supreme Court publishes essentially all of its final decisions.⁵⁰ The corpus therefore approximates the full population. A comparison with the court's annual statistics confirms its coverage.⁵¹ Across 2011-2025, the corpus captures 97.8% of all civil-appeal dispositions reported in the Geschäftsberichte. It contains 24'617 of 25'161 dispositions. Annual coverage ranges from 97.2% to 98.4% (Table 2 and Figure 2).

[33] The residual of about 2% remains unattributed. Publication or capture lag is a plausible cause. The residual is not broken down further here. Given this coverage, the publication-selection caveat conditioning Parts I and II does not arise in the same form.⁵²

3.1. Sample Construction

Stage	4A (First Civil Law Division)	5A (Second Civil Law Division)	Total
Raw judgments 2011-2025	9'796	14'821	24'617
Usable full text (>5'000 char.)	7'702	10'127	17'829
Pre-selection frame	5'883	3'543	9'426
Pecuniary merits judgments	4'628	1'947	6'575
of which with an available value in dispute	1'679	328	2'007
of which within the value window	1'351	279	1'630
Pecuniary non-entry decisions within the value window	130	46	176
Analysis scope (merits and non-entry)	1'481	325	1'806
less R-44 (purely procedural federal object)	1'143	217	1'360
less R-42 and R-45 (first-instance procedure, protection matters)	1'138	201	1'339
Cost-deviation sample (merits only)	1'138	201	1'339
Golden set (manually coded)	73	57	130

Value window: CHF 30'000 to 500 million. The golden set is drawn from the pre-selection frame. *Source:* swiss-caselaw database; own compilation.

Table 3: Corpus construction and golden set, by civil-law division.

49 KÖNIG, Part I (Fn. 5), paras. 32 and 190; KÖNIG, Part II (Fn. 5), paras. 31 and 108 f.

50 Art. 27 FSCA charges the court with informing the public about its case law, requires publication as a rule in anonymised form and delegates the governing principles to a regulation (Art. 27(3) FSCA). Art. 59(1)(a) and (b) of the Federal Supreme Court Regulations (FSCR) of 20 November 2006 (SR 173.110.131) prescribes online publication of all decisions included in the official reports and of all final and partial decisions. Full publication of the final decisions is therefore prescribed, while the official reports (BGE) carry only a curated selection of leading cases. "Essentially all" accordingly means the full set of final decisions, less the small publication or capture lag quantified below.

51 The annual reports (Geschäftsberichte) are published by the Federal Supreme Court at <https://www.bger.ch/index/federal/federal-inherit-template/federal-publikationen/federal-pub-geschaeftsbericht.htm>.

52 For that caveat, see KÖNIG, Part I (Fn. 5), paras. 32 and 190; KÖNIG, Part II (Fn. 5), paras. 31 and 108 f. The disposed-appeal counts are taken from the rows "Beschwerden in Zivilsachen" for the two civil-law divisions, specifically the "Erledigt" column in each report's statistical annex. For the primary sources, see Swiss Federal Supreme Court, Geschäftsberichte des Bundesgerichts (Federal Supreme Court Annual Reports): 2011, p. 28; 2012, p. 26; 2013, p. 24; 2014, p. 26; 2015, p. 26; 2016, p. 22; 2017, p. 22; 2018, p. 24; 2019, p. 24; 2020, p. 24; 2021, p. 24; 2022, p. 24; 2023, p. 26; 2024, p. 24; 2025, p. 32.

[34] The court assigns docket prefixes to civil appeals in its two civil-law divisions. These are 4A for the First Civil Law Division and 5A for the Second Civil Law Division. Together, the prefixes identify the appeals in civil matters studied here (Beschwerde in Zivilsachen, Section 2.1). The subsidiary constitutional complaint (4D/5D) is a reduced remedy confined to constitutional grievances where no ordinary federal appeal is available. In this civil docket, it principally concerns cases below the value-in-dispute threshold. It is excluded, mirroring the appeal-only delimitation of Part II.⁵³ After deduplication, this yields the 24'617 judgments.

[35] Restricting the corpus to usable full texts of more than 5'000 characters removes short procedural dispositions with little cost content. Examples include inadmissibility and write-off decisions, single-judge orders and three-member-panel decisions (Section 2.3). The restriction leaves 17'829 judgments. A proceeding-type pre-filter then removes summary matters, provisional measures, debt-enforcement supervision, arbitration, and child and adult protection. This yields the pre-selection frame of 9'426 judgments. The restriction is a design delimitation for procedural and cost comparability. It does not claim that every excluded category lies outside the court's full legal cognition. Table 3 reports the funnel.

[36] Within this frame, coded criteria delimit the analysis scope. They are applied on reading at the extraction stage. The matter is pecuniary and a value in dispute is available. In most cases, the court itself states a figure.⁵⁴ A uniform value window from CHF 30'000 to CHF 500 million then yields the analysis scope of 1'806 appeals. The floor mirrors the ordinary admissibility threshold (Section 2.1) but remains a study-design restriction. Appeals from a sole cantonal instance or the Federal Patent Court may be admissible below it (Section 2.1). The chosen floor is deliberately stricter than the CHF 15'000 threshold for tenancy and employment matters. Employment disputes up to and including CHF 30'000 carry a flat fee assessed independently of value (Section 2.2). Because that fee is not value-based, these disputes lack a value-based court-fee baseline. Accordingly, employment appeals at exactly CHF 30'000 are excluded from the court-fee deviation analysis. The cap guards against extreme or implausibly stated values. Robustness checks lower the floor to CHF 15'000 and trim the highest stakes (Section 5.7).

[37] The cost-deviation analysis additionally restricts the scope to the 1'630 judgments decided on the merits. Three adjudicated exclusions then leave 1'339 cost-deviation cases.⁵⁵

[38] The appeal-outcome and second-cost-level analyses retain the full scope because the disposition is precisely what they study. The statutory value-in-dispute regime (Section 2.1) structurally characterises the population as high-value. Three exclusions apply to the cost-deviation sample.

53 KÖNIG, Part II (Fn. 5), para. 34.

54 In some judgments, the extraction reconstructs the current value from the claims recited in the text. Its reasoning notes that the judgment states no separate figure. A text classifier over that reasoning flags 238 of the 1'339 deviation cases. The flag is a mechanical text diagnostic.

55 A merits judgment is a substantive judgment, as opposed to an inadmissibility, write-off or settlement disposition. The non-merits dispositions follow a distinct procedural path. Party compensation may be reduced where a proceeding ends without a merits judgment (Section 2.2). The exclusion is a design delimitation between merits and non-merits decisions. It does not assume that every non-merits court fee lies at the bottom of its statutory band.

Each is adjudicated case by case on the judgment's full text and recorded in a machine-readable layer.⁵⁶

[39] R-44 removes appeals adjudicating exclusively a procedural question because slight procedural effort can affect both cost components. This confounds the tariff relationship measured here.⁵⁷ This rule removes 270 of the merits judgments and leaves 1'360. Across the full scope of 1'806 appeals, R-42 flags 54 appeals on reading as first-instance proceedings that were not full ordinary ones. R-45 flags 10 appeals in the child and adult protection category of the proceeding-type pre-filter. That category expressly includes adoption.⁵⁸

[40] Within the merits judgments, R-42 flags 31 cases and R-45 flags 7. Their union contains 37 merits judgments. R-44 already removes 16 of those judgments. R-42 and R-45 therefore remove another 21 cases and leave the 1'339 judgments. The exclusion layer does not rewrite the coded first-instance proceeding type. It records the adjudicated exclusion separately. The coded field and the selection decision thus remain distinguishable.

[41] All three exclusions act on the cost-deviation sample alone. The appeal-outcome and second-cost-level analyses keep the full scope of 1'806 because a decision not to enter the matter is itself an outcome. Removing these cases would strip 166 of 176 non-entry decisions from the analysis that studies them. The trilingual frame comprises 4'852 German, 3'790 French, and 784 Italian judgments. It divides into 5'883 appeals in the First Civil Law Division (4A) and 3'543 in the Second Civil Law Division (5A). The Second Civil Law Division contributes proportionally fewer full-cognition appeals because much of its caseload is removed by the pre-filter. This includes family-law summary matters and debt-enforcement cases.

[42] The extracted lower-instance court is the authority whose decision is contested. Its name does not by itself identify whether it acted on appeal or as a sole instance. A blind full-text scan of all 1'339 cost-deviation cases identified at least 77 supplementary-insurance appeals from a sole

56 The residual error of that adjudication is measured. A stratified random sample of 80 already adjudicated appeals was re-read blind by a second reader who did not see the first verdict: 50 drawn from those classified as not purely procedural and 30 from those classified as purely procedural. The two readings agree on 78 of the 80. For the contamination risk of leaving a purely procedural appeal in the sample, they agree in all 50 cases. This places the one-sided 95 per cent upper bound on that error rate at 5.8 per cent. For the opposite risk of wrongly removing a merits appeal, they disagree in two of 30 cases. The upper bound is 19.5 per cent. Both were resolved in favour of the first verdict on joint re-reading, but that resolution was not blind and the blind figure is the one reported here. The measurement covers the 910 confident verdicts available when the sample was drawn.

57 Each component has its own authority and each authority concerns a single case. For the court fee, BGer, 5A_460/2018, 11 June 2018, consid. 7: on account of the slight effort incurred, the Court declined to exhaust the tariff range and set the fee at CHF 3'000, observing that the value in dispute of over CHF 15 million would readily carry up to CHF 100'000 under Art. 65 FSCA and the FSFT. For party compensation, BGer, 4A_546/2013, 13 March 2014, consid. 5: in view of the slight effort required by the merely procedural questions, the Court set compensation below the minimum rate under Art. 4 FSPCR by applying Art. 8(2) FSPCR. No court fee was imposed there because the losing party was a canton (Art. 66(4) FSCA). Neither passage states that the tariff minimum is undercut as a rule.

58 The completeness of that third set is measured. Excluding only the cases already known to be affected would repeat the oversight that the rule guards against. A counter-search therefore covered the full text of all 1'806 scope judgments. It used the protection and child-abduction vocabulary of the pre-filter itself, together with the names of the protection authorities and the corresponding provisions of the Swiss Civil Code (SCC) of 10 December 1907 (SR 210). The search was required to recover all ten known cases before its output was read and it returned 126 judgments. The 116 not already excluded were each adjudicated on the full text and each retained. In every one the governing subject matter is an ordinary civil-law question and the term that triggered the hit is incidental to it. Counting each axis once per judgment, the trigger was the name of a protection authority in 52 of the 116, the protection vocabulary in 50, the child-abduction vocabulary in 21 and a cited provision of the SCC in 4. Ten judgments were reached on more than one axis. One of the ten is an adoption matter, retained in the set precisely because the bound pre-filter category names adoption. Reading the category more narrowly would have re-introduced the very category shift the rule is meant to prevent. The exclusion set is therefore complete within this search. Its reach is the scope population.

cantonal instance (5.75 per cent). This is a functional lower-bound classification. An insurance-court name located 67 cases. An express sole-instance statement identified eight additional cases despite another court label. A direct action apparent from the procedural history identified the remaining two. The date of the contested cantonal judgment is recorded for all 77. In two cases, that judgment predates the CPC's entry into force on 1 January 2011. Neither case, however, establishes a pre-CPC basis for the sole cantonal instance. In one case, the Federal Supreme Court expressly grounded that jurisdiction in Art. 7 CPC and the cantonal implementing act in force from 1 January 2011. It also applied the CPC to the appeal because Art. 405(1) CPC turns on the judgment's dispatch date. The other judgment made no finding on the point.⁵⁹

[43] Three further appeals in the cost-deviation sample came directly from the Federal Patent Court. The first-instance value in dispute is retained only as auxiliary metadata. It is not substituted for the amount used to construct the federal tariff baselines. Independently appealable preliminary and interim decisions (Section 2.1) can enter when the remaining criteria are met.⁶⁰

3.2. Coding and Extraction

[44] A random sample of 130 judgments was manually coded as the ground-truth golden set.⁶¹ It comprises 73 German, 46 French and 11 Italian judgments. Of these, 73 are from the First Civil Law Division and 57 from the Second Civil Law Division. Coding follows a fixed codebook of 63 variables. The codebook extends the schemas used in Parts I and II with effort indicators specific to the federal appeal.⁶² The added indicators cover panel size, the exchange of submissions, suspensive effect, joined parallel appeals and the question of fundamental importance.⁶³ The golden set is the reference for calibrating and validating the LLM-assisted extraction. Both steps precede application to the full frame. This sequence follows the prompt-engineering approach of Parts I and II.⁶⁴

59 BGer, 4A_118/2011, 11 October 2011, consid. 1.1-1.3 (cantonal judgment of 7 December 2010, dispatched on 12 January 2011; sole cantonal instance under Art. 7 CPC in conjunction with § 14 of the Aargau Introductory Act to the Swiss Civil Procedure Code (CPCIA-AG) of 23 March 2010 (SAR 221.200), in force since 1 January 2011; the implementation period granted to the cantons under Art. 130(2) FSCA for the double-instance requirement of Art. 75(2) FSCA expired with the entry into force of the CPC). BGer, 4A_5/2011, 24 March 2011, is cited for the date of the cantonal judgment (28 October 2010) alone. Its consid. 1.1 treats the Vaud Cour des assurances sociales as a superior court ruling as last cantonal instance and contains no transitional finding. The case is therefore an observed instance.

60 A separate legal review of the 182 court-fee observations whose value basis was individually examined identified 18 challenges to preliminary or interim decisions and 62 challenges to final decisions. In 102 observations the decision type was not determinable from the judgment text. The 18 are therefore a directly observed minimum.

61 Drawn with seed 42 and reproducible on the fixed corpus snapshot.

62 KÖNIG, Part I (Fn. 5), paras. 37-42; KÖNIG, Part II (Fn. 5), paras. 40-46.

63 These indicators mirror the stages of the appeal proceeding set out in Section 2.3. The full list of the 63 variables is in Appendix 1.

64 KÖNIG, Part I (Fn. 5), paras. 37-42; KÖNIG, Part II (Fn. 5), paras. 40-46.

Stage	Judgments	Field accuracy
Calibration (ranks 0001-0030)	30	0.928
Held-out validation (ranks 0031-0130)	100	0.938
of which: appeal value in dispute	100	0.770
Full-population rollout (frame)	9'426	n.a.

Field accuracy: share of scored fields reproduced exactly. For the appeal value in dispute, precision is 0.83 where the judgment states a figure, while recall is 0.50 on the wider codebook construct. *n.a.*: no ground truth for the full-population rollout. *Source*: Own calculations.

Table 4: Extraction calibration, held-out validation, and full-population rollout.

[45] The extraction reuses Part II's modular, zone-based pipeline.⁶⁵ It combines a shared 63-variable Pydantic schema with a system prompt. Each extraction target has one rule-bearing prompt. A thin layer supplies the configuration. The principal adaptation is systematic trilinguality across a nationwide corpus. The pipeline adds language detection and per-language keyword sets. These sets cover text zones, dispositive anchors and effort markers.⁶⁶ The model is OpenAI's gpt-5.4-mini. It uses structured outputs and target-specific reasoning effort. Calibration used the first 30 golden-set judgments over successive refinement cycles. A plateau test extended development from seven to fifteen rounds. Together, these steps raised clean field accuracy on the calibration set to 0.928 (Table 4).⁶⁷

[46] Following Part II, the pipeline was frozen after calibration.⁶⁸ It was then applied to 100 held-out judgments. Their texts were not processed during the fifteen development cycles. Out-of-sample field accuracy is 0.938, marginally above the calibration figure. This indicates that the refined rules generalise beyond the development sample.⁶⁹ The headline figure averages over the whole schedule. The field anchoring the deviation model is weaker than that average. The appeal value in dispute is reproduced exactly in 0.77 of the held-out judgments. Where the court writes a figure, the extraction reproduces it reliably, with a precision of 0.83. But it records a value in only half the judgments for which the coding rules would derive one. The field therefore errs predominantly by omission, though not exclusively. Four of its twenty-four positive held-out predictions are spurious or wrong in amount.⁷⁰

65 Part II sets out that extraction pipeline, its calibration on a manually coded reference set and its held-out validation in full (KÖNIG, Part II (Fn. 5), paras. 40-46). The present paper reuses the architecture with a trilingual configuration and the Swiss federal-appeal schema and reports the resulting data.

66 The keyword sets were verified across the three languages, for example "Demnach erkennt das Bundesgericht" / "Par ces motifs, le Tribunal fédéral prononce" / "Per questi motivi, il Tribunale federale pronuncia" for the dispositive anchor.

67 The accuracy figures score the analytical variables. Two coder-side quality fields, the case identifiers and the free-text dispute-value reasoning lie outside the scored set as annotation or identification fields. The value in dispute and all cost variables are fully scored.

68 KÖNIG, Part II (Fn. 5), paras. 40 and 42.

69 Three documented adjustments fall between the close of calibration and the held-out run. All changed the pipeline to match the ground truth. None changed the ground truth to match the pipeline. Two implement refined coding rules and were measured on the calibration cases alone. The third aligns the schema's spelling of the partial-grant category with the coder's short form. A final plausibility check over all 130 coded cases revealed this discrepancy in six held-out cases. It changes how the category is written. Without it the held-out accuracy is 0.937, not 0.938.

70 The consequences for the sample are taken up in Section 7. The positive side was additionally audited by hand on a stratified risk sample of forty extractions (band edges, values above CHF 10 million, threshold echoes at CHF 30'000, flagged derivations). Thirty-five values were court-stated or correctly derived, two misread values were corrected through a documented correction layer applied at panel construction and three collocation or cost-only appeals record the underlying claim where the case law would measure the value differently, without moving the applicable court-fee band. They are retained and disclosed.

[47] Procedure duration is not recorded in the database. It was recovered from the filing date stated in the judgment text as a separate extraction target. Against the golden set, 81 of 84 extracted durations were within ± 3 days. Duration is reported as a robustness covariate and kept out of the exact-match headline. The calibrated configuration was then rolled out over the full frame of 9'426 judgments (Section 3.1) via the batch API. At an extraction cost of approximately USD 221, this produced the population dataset used in the analysis.⁷¹

4. Methodology

[48] This section describes the variables, tariff-baseline computation and regression design. These elements relate court-awarded litigation costs to the statutory federal tariff and observable case and procedural characteristics. The design adapts the series' deviation-from-tariff approach to the Federal Supreme Court.⁷² Two features simplify that design. A single federal tariff applies, so court-level heterogeneity needs no modelling, and one court sets both cost components, so judge effects vary only within it.

4.1. Variables

[49] The variables capture general case information, relevant cost components, and proxies for procedural effort. Their definitions reflect both the regression design and the information available in published judgments. Appendix 1 lists them in full.

4.1.1. Dependent Variables and Tariff Baselines

[50] The primary outcomes are the court fee and party compensation (Section 2.2). Both are set in the appeal judgment's operative part and measured in nominal Swiss francs. As in Part II, the data also record the lower-instance costs reviewed in each judgment (Section 2.2).⁷³ Because the reproduced amounts are the noisiest extracted field, the secondary analysis models only whether the lower-instance costs were modified.

[51] For each judgment, the study computes a mechanically implied baseline for each cost component. Each baseline uses the relevant federal schedule and the cost-relevant value in dispute before the Federal Supreme Court (Section 2.2). The court-fee baseline is the midpoint of the applicable FSFT band for pecuniary disputes. The party-compensation baseline is the midpoint of the applicable value-based FSPCR band (Section 2.2). Formally, for each judgment i ,

$$\text{Baseline}_i^k = f^k(V_i), \quad k \in \{CF, PC\}, \quad (1)$$

where V_i is the cost-relevant value in dispute before the Federal Supreme Court (Section 2.2).

71 The residual extraction errors concentrate in inherently interpretive fields. The court-awarded amounts, which are the dependent variables, are extracted reliably out of sample. The value in dispute, by contrast, is among the least reliably extracted fields and is stated as a figure in only about a quarter of all judgments, so the deviation regressions run on the judgments with an available value (Section 7).

72 KÖNIG, Part I (Fn. 5), para. 72; KÖNIG, Part II (Fn. 5), para. 59.

73 KÖNIG, Part II (Fn. 5), para. 65.

[52] The functions f^{CF} and f^{PC} map that value to the applicable band midpoint. They implement the FSFT and the FSPCR, respectively. Both schedules remained unchanged over the 2011-2025 window. The functions are step functions of the value in dispute. They remain constant within a band and rise at the band edges. The compensation scale has a single linear segment above CHF 5 million.⁷⁴ Because awarded party compensation includes value-added tax (Section 2.2) while the tariff scale is net, the award is VAT-normalised before comparison.

[53] The applicable VAT rate follows the judgment year. The recipient's Swiss domicile serves as the VAT indicator. Over the observed range, the midpoint's share of the value in dispute declines overall, despite upward jumps at all nine band boundaries. The measured baseline elasticity captures the schedule's degressive pattern over that range. Using the baseline therefore nets out its less-than-proportional rise. The log-baseline coefficient measures the award's remaining algebraic deviation after accounting for that structure.

[54] The band midpoint is an analytical anchor, not an amount the court owes. Within the statutory range, the value in dispute is only one of the several criteria that fix the court fee, and the Court retains discretion in setting it (Section 2.2). The court may fix the fee by reference to all of them even when the cost-relevant amount is unknown. The midpoint thus need not be an amount from which the court consciously departs.

[55] Party compensation is fixed by the case's importance and difficulty. The required work and time also matter (Section 2.2).⁷⁵ An award below the midpoint therefore lies in the lower half of a lawful range and is not a departure from the tariff. The midpoint is used because it is the one point in each band fixed by the schedule alone. It is available for every judgment without discretionary input and thus provides a reproducible benchmark. Section 5.7 re-anchors the estimates on the band minimum and maximum to show that the slope does not depend on this choice.

4.1.2. Case Characteristics

[56] Case characteristics record factual features of the dispute and parties. They include the number of parties and attorneys on each side. For the lead appellant and respondent, they also include Swiss domicile, legal form, and industry (NOGA) classification. The case caption's legal subject is grouped by keyword into broad civil-law categories.⁷⁶

4.1.3. Court-Side Characteristics

[57] Court-side characteristics identify the panel's judges and court clerk. The presiding judge is flagged separately to allow judge fixed effects. Panel size is the primary proxy for case difficulty and effort. Panel-size categories enter as dummies.⁷⁷

74 Above that threshold the compensation scale is itself a range, from CHF 20'000 to one per cent of the value in dispute (Section 2.2), so its midpoint rises linearly with the value. The segment covers 34 observations.

75 For an illustration, BGer, 5A_237/2025, 19 January 2026, consid. 4, setting compensation at CHF 2'000 within the applicable range of CHF 1'500 to 6'000 by reference to the value in dispute and the difficulty of the matter.

76 Contract, family, employment, tenancy, property, succession, and so on. The full grouping is in Appendix 1.

77 Three judges as the rule, five for questions of fundamental importance or at a judge's request and a single judge in the simplified procedure (Section 2.3).

4.1.4. Cost Allocation

[58] Cost allocation follows the loser-pays principle (Section 2.2). The appellant's loss share captures the proportion of the appeal that was unsuccessful. A value of 0 denotes full success and 1 denotes full defeat. For party compensation, this realised ex-post allocation control enters only after the procedural block. Earlier models show the baseline slope before conditioning on outcome.

4.1.5. Procedural Complexity Indicators

[59] The procedural indicators are objective proxies for effort. They are read from the procedural history along the stages of Section 2.3. They record whether an exchange of submissions was invited. They also record whether a second exchange or further unsolicited submissions followed. Further indicators cover a public hearing or deliberation.

[60] They record decisions on suspensive effect or other provisional measures. Where a request was made, they record whether it was granted or dismissed. They also cover joined parallel appeals, proceeding duration, and decision length. Duration and decision length are realised ex post, jointly with the court-awarded amounts. Specifications conditioning on them therefore estimate compression net of realised effort. The ex-ante prediction model excludes them and all other ex-post information (Section 6).

4.2. Regression Design

[61] The court fee and party compensation are estimated separately, as they follow different schedules. The main specification is a deviation model. Its dependent variable is the log ratio of the court-awarded amount for the relevant component to its implied tariff baseline. The sample comprises pecuniary judgments decided on the merits with an available value in dispute.

[62] The full specification is

$$\ln\left(\frac{Award_i}{Baseline_i}\right) = \beta \ln Baseline_i + \gamma' X_i + \lambda L_i + \alpha_t + \mu_j + \varepsilon_i, \quad (2)$$

where β is the log-baseline coefficient. The vector X_i collects case and procedural controls with coefficients γ . The term L_i is the appellant's loss share with coefficient λ . It enters only for party compensation and is realised ex post. The terms α_t and μ_j are year and presiding-judge fixed effects. The term ε_i is the error.

[63] Because the log baseline enters the dependent variable through the denominator, the award's implied level elasticity with respect to the tariff is $1 + \beta$. The log-log form is motivated by the right-skewness of the cost components and by the schedules' degressive structure over the observed range. A zero log-baseline coefficient indicates a proportional tariff anchor, while a negative coefficient indicates stronger downward deviation at higher values in dispute. Estimation proceeds through nested models. Year fixed effects enter first (M2).

[64] Case characteristics follow (M3), including the type of first-instance proceeding.⁷⁸ Procedural-complexity indicators enter next (M4). This block includes panel-size dummies, log duration, and log decision length, each with a missing-value flag. For party compensation, the appellant's loss share then enters as the cost-allocation term (M5). Presiding-judge fixed effects complete the sequence (M6). Because a single federal court applies a single tariff, the court dummies and canton-by-baseline interactions used in Part II are unnecessary.⁷⁹ Models for a given cost component use one sample that is complete on all M6 covariates. The resulting M6 samples contain 1'315 court-fee and 973 party-compensation observations. This keeps their incremental explanatory contributions comparable. Inference reports heteroskedasticity-robust (HC1) standard errors throughout.

[65] For the headline coefficient, inference uses a wild-cluster bootstrap by presiding judge. The presiding judge is the clustering level linked to the decision-making unit. The period has only about ten distinct presiding judges and an effective count near four. The wild-cluster bootstrap is therefore preferred to the cluster-robust sandwich. It draws Webb six-point weights. These are recommended with so few clusters.⁸⁰ Part II uses them for the same reason.⁸¹ By contrast, two-point Rademacher weights span only 2^G sign vectors at G clusters. They cannot resolve a p -value below $2/2^G$. The cluster-robust sandwich's unsuitability is measured.⁸²

[66] The bootstrap draws 99'999 replications with a fixed seed. A reported p -value near 0.01 therefore carries a simulation standard error of about 0.0003. Two further seeds move the same specification by 0.0004 for the court fee and 0.0007 for party compensation. For a single run, 0.0002 and 0.0004 are predicted by the binomial expression $\sqrt{p(1-p)/B}$. The second reported digit is therefore stable. The third is not.

[67] Confidence intervals for the level elasticity use test inversion, not a bootstrapped standard error. The confidence set contains null values that the same wild-cluster bootstrap does not reject at five per cent. The interval thus inherits the test's few-cluster correction. It need not be symmetric around the point estimate. Inversion is first evaluated on a grid of null values. Bisection then refines each bound, so grid spacing does not limit resolution.

[68] The inversion uses 19'999 replications per null value. After bisection, a bound's Monte Carlo uncertainty is on the order of 0.001 to 0.002. This uncertainty exceeds that of a p -value because a

78 Three dummies for a simplified, a special, and an undeterminable first-instance proceeding, with the ordinary proceeding as the reference. A simplified proceeding under Art. 243 CPC and an ordinary one differ in form and in effort, so leaving that axis out would leave a plausible confounder unobserved. The four appeals that reached the court by an administrative route run in the reference category. Four observations would not support an informative separate indicator. The specification is reported both with and without these controls (Appendix 2, R17). The decision to add them preceded inspection of the estimates. Reporting both specifications makes the choice visible and avoids selection among researcher degrees of freedom.

79 KÖNIG, Part II (Fn. 5), para. 62.

80 ADRIAN C. CAMERON/JONAH B. GELBACH/DOUGLAS L. MILLER, Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics* 90.3 (2008), 414-427, secs. 1 and 3.2-3.3; JAMES G. MACKINNON/MATTHEW D. WEBB, The Wild Bootstrap for Few (Treated) Clusters. *The Econometrics Journal* 21.2 (2018), 114-135, secs. 2.2 and 6.

81 KÖNIG, Part II (Fn. 5), para. 114.

82 A simulation on the design of the earlier specification (the same ten presiding-judge clusters, $n = 1'574$ court-fee and 1'126 party-compensation observations, $S = 1'000$ replications of $B = 399$ each, with a presiding-judge-specific slope imposed as the null) rejects at 11.7 and 11.2 per cent where the nominal level is five. The realised coverage of the cluster-robust t interval is therefore 88.3 and 88.8 per cent, not the nominal 95. The figures are stated with their reach. They were measured on the specification that preceded the first-instance proceeding-type controls and the coverage of the present M6 specification has not been simulated. Only the direction of the finding is used here. The paper therefore reports the bootstrap throughout instead of the sandwich.

bound locates where the p-function crosses the level instead of evaluating it at a fixed point. The third decimal of an interval bound should accordingly be read as an indication.

[69] The implied level elasticity $1 + \hat{\beta}$ accompanies each coefficient. It is the deviation model's natural interpretation. It equals the slope from directly regressing the log award on the log baseline. The proportional benchmark is one. A value below one is compression. Least squares in logarithms identifies the conditional mean of the log award. Exponentiation yields the conditional geometric mean. The term typical award is used in that sense throughout.

[70] A quantile regression separately estimates the conditional median as a robustness check (Section 5.7). The Poisson specification below targets the conditional mean in levels. The dependent variable contains the baseline as $-\ln \text{Baseline}$, while the regressor contains it as $+\ln \text{Baseline}$. Because the anchor appears on both sides, classical measurement error would attenuate this elasticity towards zero. Equivalently, it biases the coefficient towards -1 . Model-free descriptive evidence corroborates the regression. It comprises the binned median award-to-baseline ratios in Figure 4 and the quartile medians in Section 5.1. The same less-than-proportional pattern appears without relying on the deviation specification.

[71] The two cost components are handled differently because party compensation is frequently not awarded at all. A zero is a deliberate non-award, or corner solution. The court fee is almost always positive and is estimated on positive awards. Party compensation uses a two-part hurdle model. The first part is a logit for whether any positive award is made,

$$\Pr(PC_i > 0 \mid W_i) = \Lambda(\omega'W_i), \quad (3)$$

where Λ is the logistic function. The vector W_i contains participation covariates, chiefly whether a response was invited. The vector ω contains the associated coefficients.

[72] Conditional on a positive award, the second part is a positive-only deviation regression. It follows (2) and is directly comparable to the court-fee specification. A Tobit model is deliberately not used because the zeros are corner solutions and not left-censoring. Instead, a Poisson pseudo-maximum-likelihood estimator operates on the cost level, $\mathbb{E}[\text{Award}_i \mid x_i] = \exp(x_i'\theta)$. Here x_i contains the model covariates and θ the associated coefficient vector. The estimator admits zeros and provides the zero-inclusive robustness check on elasticity.⁸³

5. Results


[73] The cost-deviation analysis rests on 1'339 pecuniary merits judgments with an available value in dispute (Section 3.1). These judgments survive the three adjudicated exclusions in Section 3.1. Of the 1'333 judgments with positive court fees, 1'315 enter the court-fee deviation sample (Section 4.2). The other 18 are employment appeals at exactly CHF 30'000 and carry no value-based fee baseline. A total of 973 judgments show positive party compensation (Section 4.2). The appeal-outcome and second-cost-level analyses use the broader scope of 1'806 pecuniary appeals with an available value in dispute (Section 3.1).

83 JOÃO M. C. SANTOS SILVA/SILVANA TENREYRO, The Log of Gravity. Review of Economics and Statistics 88.4 (2006), 641-658, secs. 2.1 and 3.1.

[74] This scope additionally includes the 176 non-entry decisions and is not narrowed by the three exclusions (Section 3.1). The deviation sample excludes them because they follow a distinct non-merits procedural path (Section 3.1). For party compensation in non-merits proceedings, the tariff prescribes a reduction (Section 2.2). The court fee has no automatic reduction. It is assessed within the statutory range (Section 2.2).⁸⁴

[75] The median value in dispute is CHF 130'000 in the cost-deviation sample and CHF 125'000 across the analysis scope. The cost-deviation sample is trilingual, comprising 742 German, 466 French, and 131 Italian judgments. Nearly all these cases are decided by three- or five-judge panels.⁸⁵ Results proceed from the shape of court-awarded litigation costs to tariff compression in the deviation model. They then cover procedural and party-compensation margins, appeal outcomes, legal subjects, and lower-instance cost review.

5.1. The Shape of Court-Awarded Litigation Costs



Component-specific positive-award samples drawn from the cost-deviation sample. *Source:* Own calculations.

Figure 3: Distribution of the ratio of the court-awarded amount to the tariff baseline, by cost component.

⁸⁴ The outcome tabulation of Table 6 runs over the 1'804 appeals of this scope with a determinable outcome. Two appeals have no determinable outcome. The outcome and the disposition type are coded independently and agree on the 176 non-entry cases.

⁸⁵ The sample contains 1'157 three-judge and 181 five-judge panels. One further appeal was decided by a two-member bench, as its case caption confirms. It is folded into the three-judge reference category for the estimation. No appeal in the cost-deviation sample was decided by a single judge.

[76] Court-awarded litigation costs are modest relative to the high values in dispute and sit in the lower half of the tariff band. The median awarded court fee is CHF 5'000, with an interquartile range of 2'500-6'500. Median VAT-normalised party compensation is CHF 5'555, with an interquartile range of 2'777-7'400. Relative to the mechanically implied tariff midpoint, the median award is 0.83 for the court fee and 0.62 for party compensation. Thus, the typical award lies a sixth to nearly two fifths below the statutory band's midpoint while remaining inside it (Figure 3).


[77] The award-to-midpoint ratio is concentrated below parity for both components. Court-fee awards cluster at and just below the midpoint. Party compensation forms a single mode well beneath it. The gap to the midpoint is uneven across the value in dispute. Across value-in-dispute quartiles, the median party-compensation ratio stays near 0.6, while the court-fee ratio starts at about two thirds in the lowest quartile. It reaches parity in the lower-middle range and about seven tenths in the top quartile. This is a first descriptive sign of the less-than-proportional pattern quantified in Section 5.2.

5.2. Tariff Compression

DV: ln(award/baseline)	Court fee				Party compensation			
	M1	M3	M4	M6	M1	M3	M4	M6
ln(baseline)	-0.104*** (0.020)	-0.180*** (0.020)	-0.211*** (0.020)	-0.213*** (0.020)	-0.029 (0.032)	-0.126*** (0.029)	-0.122*** (0.031)	-0.122*** (0.029)
implied elasticity	0.90	0.82	0.79	0.79	0.97	0.87	0.88	0.88
response invited (Art. 102 FSCA)			0.074** (0.032)	0.082** (0.035)			0.223* (0.119)	0.240** (0.120)
appellant loss share								0.009 (0.040)
Year FE	no	yes	yes	yes	no	yes	yes	yes
Case characteristics	no	yes	yes	yes	no	yes	yes	yes
Procedural complexity	no	no	yes	yes	no	no	yes	yes
Judge FE (presiding)	no	no	no	yes	no	no	no	yes
N	1'315	1'315	1'315	1'315	973	973	973	973
adj. R ²	0.031	0.151	0.187	0.191	0.001	0.075	0.111	0.139

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. HC1 robust standard errors in parentheses. OLS, positive awards only. The specifications are nested. M2 and M5 are omitted. FE: fixed effects. Bold: full model (M6). Source: Own calculations.

Table 5: Tariff compression: deviation of court-awarded litigation costs from the statutory federal tariff baseline.



Binned medians, positive awards. *Source:* Own calculations.

Figure 4: Court-awarded litigation costs versus the federal tariff baseline.

[78] The deviation regressions confirm the pattern and quantify it. Table 5 reports the nested models for both cost components. The log-baseline coefficient is negative throughout and highly significant in the controlled specifications. In the fully controlled specification (M6), it is -0.213 for the court fee and -0.122 for party compensation. The corresponding level elasticities are 0.79 and 0.88, respectively. Both are below the proportional value of one. Court-awarded litigation costs therefore rise less than proportionally with the tariff itself. Because each tariff anchor itself rises less than proportionally with the value in dispute in the relevant component-specific estimation sample, the elasticities with respect to the stakes are lower still.

[79] Replacing the log baseline with the log value in dispute gives 0.42 for the court fee and 0.43 for party compensation.⁸⁶ Compression is significant under the presiding-judge wild-cluster bootstrap using Webb six-point weights. These weights are appropriate for so few clusters. The p-values are 0.0050 for the court fee and 0.0152 for party compensation. Figure 4 shows the same fact descriptively. Binned median awards lie below the proportional line. The court-fee series bends further below it as the tariff rises, while party compensation remains below the proportional line throughout at a roughly constant fraction of the tariff.

⁸⁶ The two readings are linked by an algebraic identity on their respective component-specific estimation samples (Appendix Table 33). The baseline rises with an elasticity of 0.48 for the court fee and 0.46 for party compensation with respect to the value in dispute (Section 2.2.3), while the award rises with an elasticity of 0.42 and 0.43. Equivalently, the log award-to-baseline ratio has a value slope of -0.06 and -0.03 . These terms describe the same covariance structure and do not identify separate statutory and judicial causal shares. The figures reported throughout and in Table 8 for the series are the elasticities with respect to the tariff baseline, which is the quantity the deviation model identifies.

5.3. Procedural Effort Within the Band

[80] Where an award falls within the tariff band is associated with procedural effort. The single strongest margin is whether a response was invited (Section 2.3). Invited-response cases carry court fees about 8.5 per cent higher (M6). Party compensation is about 27 per cent higher but is estimated less precisely. Within the merits docket, panel size is the difficulty proxy but differences by panel size are modest.

[81] Awards by five-judge benches for questions of fundamental importance differ little from those of the three-judge default. The single-judge inadmissibility track carries a near-floor fee (Section 2.3). It lies outside the cost-deviation sample by construction. The procedural block lifts the court-fee adjusted R^2 from 0.15 to 0.19. Realised effort is therefore associated with part of the within-band position, alongside the earlier case characteristics.

5.4. The Party-Compensation Margin

[82] Party compensation differs from the court fee in that it is frequently not awarded at all. Among the merits appeals it is zero in 366 of the 1'339 cases with a determinable award (about 27%; Section 3.1). The hurdle logit ties this margin closely to whether the opposing side was invited to respond. An invitation occurred in 76.6% of merits appeals. The probability of a positive award is 0.90 when a response was invited and 0.16 when it was not. The adjusted odds ratio is about 74. That association is in large part definitional.

[83] Because a respondent who is never invited will often incur no compensable effort, the invitation is closer to a precondition than to a separate cause of the award. The odds ratio measures the strength of this institutional link. Exceptions occur in both directions. Fifty of the 973 positive awards (Section 4.2) arose where no response had been invited. Conversely, 103 of the 1'026 invited proceedings ended without an award.⁸⁷ The non-award also tracks disposition. Across the broader scope, party compensation is zero in 30% of merits judgments. It is zero in 73% of non-entry decisions.


⁸⁷ In the manually coded golden set the invitation is read from the procedural history alone. The automated extraction reads it from the same history and, where that history is silent, infers it from an award of party compensation to the respondent. That inference runs from the outcome to the predictor, so it could overstate the link. The golden set, coded without it, shows the same association more sharply than the extracted population, so the pattern is not an artefact of the extraction rule.

5.5. Appeal Outcomes and Legal Subjects

	N	Share / rate
<i>A. Appeal outcome</i>		
Dismissal	1'296	71.8%
Non-entry	176	9.8%
Partial grant	170	9.4%
Full grant	125	6.9%
Remand	37	2.1%
<i>B. Appeal success rate, by value-in-dispute quartile</i>		
Q1 (low)	452	17.5%
Q2	450	17.6%
Q3	451	16.6%
Q4 (high)	451	22.0%
<i>C. Public-entity involvement (cf. Art. 66(4) and Art. 68(3) FSCA)</i>		
Public-entity party, zero court fee	141	6.4%
zero party compensation	141	58.2%
Other cases, zero court fee	1'665	1.0%
zero party compensation	1'665	32.2%

Full analysis scope, including non-entry decisions. Panel C splits by the presence of a public-entity party, which does not itself establish the statutory exemption elements. *Source:* Own calculations.

Table 6: Appeal outcomes and public-entity involvement.



Full analysis scope, subjects with at least 40 cases. Success rate from the coded appeal outcome (grant, partial grant, or remittal), not from the cost share. *Source:* Own calculations.

Figure 5: Appeal success and the zero-party-compensation share by legal subject.

[84] The recorded disposition shows that the appeal succeeds wholly or partly in only 18.4% of cases. Dismissal is the modal disposition (71.8%), ahead of non-entry (9.8%), partial grant (9.4%), full grant (6.9%) and remand (2.1%). The success rate is roughly flat across value-in-dispute quartiles and the two civil divisions (Table 6).

[85] The success rate is lower in Italian-language judgments (Table 9). Success also varies by legal subject, though within a narrow band. It ranges from about a quarter in insurance and corporate matters to about an eighth in property and contract claims (Figure 5). At a finer keyword grain, success ranges from under one in ten for purchase and works-contract claims to over one in four for family maintenance. This ordering is indicative only because the finer sub-classification is among the least reliably extracted fields (Appendix Table 50).

[86] The no-award margin varies by legal subject far more than the success rate does. The court fee is charged in almost every case.⁸⁸ Party compensation is frequently withheld and its zero share varies sharply by matter. The zero share is highest in property (76.6%), enforcement (67.4%), family (63.0%), and inheritance (60.6%). It is lowest in corporate (15.6%) and contract claims (20.0%) (Figure 5).

[87] Cases involving a public entity point towards the exemptions in Section 2.2. The extracted indicator records the entity's presence. The split is therefore an association. Party compensation is zero in 58.2% of these cases, against 32.2% otherwise (Table 6).

5.6. The Second Cost Level


[88] In each appeal, the Federal Supreme Court may also revisit the costs and party compensation of the proceedings before the lower instance (Section 2.2). This review is strongly outcome-dependent. Lower-instance cost allocation is modified in 10.7% of appeals overall. The rate is 65.6% for full grants and 60.0% for partial grants. It is 0.0% for dismissals. None of the non-entry decisions has a modification.⁸⁹

[89] The contrast is absolute on this sample. Not one of the 1'472 unsuccessful appeals carries a cost revision. Of the 332 successful appeals, 193 carry one. The table exhibits quasi-complete separation, so an unpenalised logit has no finite maximum-likelihood estimate. Appendix 2 reports the table itself with Fisher's exact test ($p = 4.8 \times 10^{-169}$ on $n = 1'804$).

⁸⁸ A zero fee occurs in up to about four per cent of the cases, depending on the legal subject. Possible statutory bases include a discretionary waiver where circumstances justify it and the public-entity exemption (Section 2.2), and a waiver after withdrawal or settlement (Art. 66(2) FSCA).

⁸⁹ A manual audit of twelve randomly drawn cases with a positive modification indicator confirms what the variable captures. In every case, the operative part addresses the costs of the proceedings before the lower instance. In eleven cases it sets them aside and remits them for reallocation. In one it resets the fee directly.

5.7. Robustness



A random sample of 350 awards per component (seed 42). *Source:* Own calculations.

Figure 6: Fitted values versus residuals, model M6.

[90] The compression result (Section 5.2) is robust across the full battery reported in Appendix 2, subject to two informative qualifications. The first concerns the estimand. A zero-inclusive Poisson pseudo-maximum-likelihood estimator operates on the cost level. It carries the same controls as the fully controlled deviation model. The conditional-mean elasticity is 0.83 for the court fee and 0.82 for party compensation. Both remain below the proportional value of one. The full control set requires 1'319 observations for the court fee and 1'339 for party compensation, including zero awards (Section 3.1). The earlier near-one estimates come from the baseline-plus-year specification. A median regression with the full control set provides a point diagnostic. Its inference is not adjusted for the few judge clusters. It returns 0.83 for the court fee and 0.95 for party compensation. Compression is therefore present at every estimand examined. It is clearest in the typical, log-scale award estimated by the deviation model. It also appears at the conditional median, more mildly for party compensation and in the zero-inclusive conditional mean in levels. The second qualification concerns the anchor. The median award lies above the band minimum but below its midpoint and maximum. At each anchor, the court-fee elasticity remains below one, at 0.86 for the band minimum, 0.79 for the midpoint and 0.76 for the maximum. Party-compensation elasticity ranges from 0.86 to 0.88 across the anchors. The direction therefore does not depend on the anchor, while the court-fee magnitude moves with it.

[91] Two further sensitivity checks leave the headline elasticities (Section 5.2) virtually unchanged. Removing the four merits appeals reaching the Federal Supreme Court by an administrative route moves the court-fee elasticity from 0.787 to 0.788. The party-compensation elasticity moves from 0.878 to 0.879.⁹⁰ Dropping legal-subject controls moves elasticity by about 0.01. VAT-normalising party compensation leaves the coefficient unchanged to two decimals. Ex-

⁹⁰ The nine administrative-route appeals divide into five direct actions and four appeals against an administrative order. The dataset carries the distinction as a descriptive flag, which is not used as a regressor because four observations would not support an informative separate indicator (Section 3.1).

cluding the functional lower-bound set of 77 supplementary-insurance cases from sole cantonal instances while retaining the M6 controls leaves headline level elasticities virtually unchanged. The court fee moves from 0.787 to 0.785 and party compensation from 0.878 to 0.882.

[92] This exclusion applies to the headline panel without component-specific value-error corrections. Those corrections are evaluated separately. Including the non-entry decisions deepens the court-fee coefficient to -0.303 . The reported merits specification is therefore conservative relative to that non-entry-inclusive estimate. Lowering the value-in-dispute threshold to CHF 15'000 moves the coefficient to -0.187 , but this variant is not a pure threshold test and supports no conservativeness claim.

[93] The added appeals below CHF 30'000 were never adjudicated under the purely-procedural rule. Thus, the threshold and selection regime change together. The confounder removed by that rule partly re-enters. The employment appeals up to CHF 30'000 stay outside the court-fee sample in either variant, because their flat fee is not value-assessed (Section 3.1). Trimming appeals above CHF 100 million or the top percentile leaves the slope essentially unchanged at 0.77. It also preserves the model-free median award-to-midpoint ratio. Compression is therefore not driven by the high-leverage tail.

[94] The results are also stable at the lower end. Judgments recording exactly CHF 30'000 may reproduce the admissibility threshold. They could bias the anchor downwards when the slope is estimated. Dropping them moves the court-fee elasticity to 0.77 and leaves the party-compensation elasticity at 0.88. Raising the floor to CHF 50'000 moves them to 0.70 and 0.91. The court-fee slope moves further from proportionality on both variants. The party-compensation slope stays well below one. Compression therefore does not rest on that artefact.

[95] The value construct is examined in the same way. The earlier extraction-reasoning text pattern marked direct figures and arithmetic sums at nearly identical rates of 17.5 and 18.2 per cent. Excluding marked rows therefore does not test value provenance. R13 instead re-estimates the fully controlled model across mechanically classified provenance groups. For the court fee, the headline elasticity is 0.787 in the court-fee estimation sample ($n = 1'315$; Sections 5.2 and 4.2). Excluding 76 appeals whose amount is not located in the judgment leaves 0.788 ($n = 1'247$).

[96] Excluding instead 143 arithmetic sums leaves 0.790 ($n = 1'184$). Excluding both leaves 0.791 among 1'116 direct figure matches. Removal counts therefore need not equal changes in the estimation sample size. A count is the provenance set removed before estimation, while n measures rows used by the component-specific complete-case estimation. Removing 76 leaves the 1'247 observations, not 1'239, because eight cases carried no court-fee observation to begin with. Each exclusion starts from the respective component-specific estimation sample.

[97] For party compensation, the full headline estimate is 0.878 ($n = 973$; Sections 5.2 and 4.2). It is 0.877 without 57 unlocated amounts ($n = 923$). It is 0.875 without 111 sums ($n = 873$) and 0.874 among 823 direct matches. All eight inverted Webb intervals remain below one. The direct-match intervals are $[0.672, 0.870]$ for the court fee and $[0.728, 0.940]$ for party compensation. These estimates describe movement across observable provenance classes. They do not identify an error-free counterfactual. An amount found in the judgment may still concern the wrong claim, procedural stage, currency, or cost component.

[98] Awards cluster on only about ten presiding judges, with an effective count near four. The wild-cluster bootstrap therefore carries the inference and its weights matter.⁹¹ Dropping the dominant presiding judge leaves the compression established in Section 5.2 significant. That judge accounts for about a third of the sample. The p-values remain 0.0173 for the court fee and 0.0121 for party compensation.

[99] The reported 95% elasticity interval is the inverted Webb interval. It contains null values that the same bootstrap does not reject. The interval runs from 0.635 to 0.872 for the court fee. It runs from 0.762 to 0.943 for party compensation. Both remain bounded below one. The cluster-robust t interval is narrower on the same data. With about ten clusters, it is anti-conservative and therefore not reported.

[100] Less-than-proportional, effort-graded compression is thus a stable feature of the Federal Supreme Court's cost-setting practice. It is clearest in the typical, log-scale award. Within the deviation model, it is somewhat sharper for the court fee than for party compensation. Model diagnostics support the specification. Fully controlled residuals are centred on zero but show greater dispersion at low fitted values (Figure 6). HC1 and the wild-cluster bootstrap address that heteroskedasticity for inference. The wide spread itself limits case-level prediction, which the ex-ante perspective quantifies (Section 6).

6. An Ex-Ante Perspective

[101] The regressions in Section 5 explain court-awarded litigation costs ex post and condition on realised effort where relevant. At filing, a party instead asks how well eventual court-awarded litigation costs can be predicted from information available at the outset. The ex-ante model therefore re-estimates the cost level using predictors known at filing,

$$\ln Award_i = \beta_0 + \beta_1 \ln Baseline_i + \gamma' Z_i + \alpha_t + \varepsilon_i, \quad (4)$$

where β_1 is the coefficient on the log baseline. Z_i collects the case characteristics known at filing, namely party characteristics and legal subject. The value in dispute enters through the baseline. The term α_t denotes decision-year fixed effects.

[102] The terms β_0 , γ , and ε_i are the intercept, control coefficients, and error as in (2). The model deliberately excludes realised-effort proxies, panel size, outcome, and presiding-judge fixed effects.⁹² Decision year lies between the ex-ante and ex-post categories. It becomes known only when the court decides, so a party cannot condition on it at filing. Removing the year effects lowers the R^2 by less than a hundredth. The ex-ante reading therefore does not depend on them.

[103] The uniform federal tariff supplies most of the ex-ante model's fit. The ex-ante model attains an R^2 of 0.76 for the court fee and 0.72 for party compensation. The tariff alone, without case characteristics, accounts for 0.71 and 0.65. Observable characteristics add the remainder.

91 At G clusters, the two-point Rademacher weights span only 2^G sign vectors. The exact p-value therefore cannot fall below $2/2^G$, or 0.002 at ten clusters and 0.004 once one is dropped. With ten clusters, exact Rademacher enumeration yields $2/2^{10} = 0.001953$. The reported Monte Carlo Rademacher value for the court fee is 0.0119 and therefore sits above that floor. The paper nevertheless relies on the Webb weights because their 6^G support removes this bound (Section 4.2).

92 KÖNIG, Part I (Fn. 5), paras. 137-140.

Part II's comparable filing-stage model attained 0.63 and 0.60 at the cantonal appellate level.⁹³ The fit here is therefore 13 and 12 percentage points higher.

[104] Part I reported a still higher first-instance fit, which its own qualification attributes to the mechanical spread of commercial-court values in dispute.⁹⁴ The tariff is therefore a useful but coarse predictor at every instance of the series.⁹⁵ A large share of the realised amount remains unexplained. Procedural effort is associated with the award's within-band position. That position is not knowable at the outset by construction. This limits ex-ante precision.

[105] A five-fold cross-validation with seed 42 returns an R^2 of 0.68 for the court fee. For party compensation, it returns 0.63. These values are 8 and 9 percentage points below the in-sample values. The cross-validation retains much of the in-sample fit. This retention shows that the in-sample fit is not solely an artefact of overfitting. The validation is internal. Its folds are drawn at random from the same fifteen years.

[106] An out-of-time test trains the model on 2011-2021 and predicts appeals from 2022-2025 that were not used in training. The resulting R^2 is 0.77 for the court fee and 0.65 for party compensation. Median forecast errors are 26% and 23%, respectively. The test runs without year effects. Their coefficients are not identified for the test period. Retaining them would make the forecast use the omitted reference year. The validation measures would then reflect that coding choice instead of the data. The R^2 figures throughout concern the log award, so they include the tariff anchor's explanatory power.

[107] The model is not confined to its estimation period. One qualification concerns levels, not fit. Party compensation in 2022-2025 is about eight per cent above the earlier-years prediction, while the court fee deviates from it by only about five per cent. Random cross-validation folds average this drift away. Because the reckoner below holds its covariates at the sample mean for the fifteen-year window, it understates current party compensation by roughly that margin.

Value in dispute (CHF)	Court fee (CHF)	Party compensation, VAT-normalised (CHF)
30'000	2'800	2'440
75'000	2'990	3'946
150'000	4'256	5'750
300'000	5'934	5'750
750'000	9'019	7'955
2'000'000	15'131	10'075

Predictions coincide within a tariff band (step-function baseline). *Source:* Own calculations.

Table 7: Ready reckoner: central predicted court-awarded litigation costs by value in dispute.

[108] Table 7 translates the model into a ready reckoner. It reports the central predicted court fee and party compensation by value in dispute. Other covariates, including the year effects, are held at their sample mean. Predicted court-awarded litigation costs rise less than proportionally with the value in dispute, reproducing the compression documented in Section 5.2. The court fee rises more than fivefold, from about CHF 2'800 to CHF 15'100, while the value in dispute rises nearly seventyfold. These figures are central predictions of the model's conditional geometric mean in the sense of Section 4.2.

93 KÖNIG, Part II (Fn. 5), para. 91.

94 KÖNIG, Part I (Fn. 5), paras. 142 f.

95 KÖNIG, Part I (Fn. 5), paras. 143-147 and 198; KÖNIG, Part II (Fn. 5), paras. 89-93 and 107.

[109] Realised award amounts scatter around the predictions. The residual standard deviation is about 0.38 log points for the court fee. For party compensation, it is about 0.47. A log-normal approximation uses this in-sample residual spread. It is not calibrated separately out of sample. The resulting ninety per cent dispersion band spans a factor of 3.5 for the court fee. For party compensation, it spans 4.7.

[110] The in-sample median absolute forecast error uses the prediction as its denominator. It is 22.7% for the court fee and 19.5% for party compensation. On an award-based denominator, the Part III errors are 22.9% and 18.1%. Part I reports approximately 21% and 26% for the commercial courts.⁹⁶

[111] The component-specific bands are tighter than Part II's factor of roughly eleven.⁹⁷ The reckoner is conditional on a positive award in the applicable component-specific model sample with an available value in dispute. For the court fee, it also requires an award set under the value-based tariff. It provides a central amount with the in-sample median absolute error. It does not model the probability of a zero award. It is therefore not a complete forecast of party-compensation risk for a typical appeal.

7. Discussion

[112] The FSFT and the FSPCR are degressive over the observed range. Their tariff anchors rise less than proportionally with the value in dispute in their respective component-specific estimation samples (Section 2.2.3). The ordinary court-fee ceiling is CHF 100'000, and the range may exceptionally extend up to twice that ceiling (Section 2.2). Proportional placement within the bands would therefore already yield court-awarded litigation costs rising less than one-for-one with the stakes. The estimates show a further negative association between the award's within-band position and the tariff baseline (Section 5.2).

[113] Against the band midpoint, the level elasticity is about 0.79 for the court fee and 0.88 for party compensation (Section 5.2). In the respective component-specific estimation samples, baseline elasticities with respect to the value in dispute are 0.48 and 0.46, while award elasticities are 0.42 and 0.43 (Section 5.2). The corresponding log award-to-baseline slopes are -0.06 and -0.03 . This is an algebraic decomposition of the same data. It does not allocate causal shares between the schedule and judicial behaviour. The residual can also reflect remaining tariff criteria, case composition, omitted variables, measurement error, selection and use of the band midpoint.

[114] The statutory criteria offer a possible reading of the observed association. The fee depends on the value in dispute, the matter's scope and difficulty, and the manner in which the proceedings were conducted (Section 2.2). Party compensation covers only the necessary costs of the prevailing party (Section 2.2). If legal effort grows more slowly than the amount at issue, awards will sit lower in the band as stakes rise. This mechanism is consistent with the data but is not tested here. The observable effort proxies capture only procedural steps, not the work they required.

⁹⁶ KÖNIG, Part I (Fn. 5), paras. 146 and 198.

⁹⁷ KÖNIG, Part II (Fn. 5), para. 93.

	Part I	Part II	Part III
Instance	first	second	third (apex)
Courts	commercial courts	cantonal high courts	Federal Supreme Court
Proceeding studied	first-instance action	appeal (Berufung)	appeal in civil matters
Tariff regimes	2 cantonal	5 cantonal	1 federal
Corpus	published judgments	published judgments	quasi-census (97.8%)
Deviation sample (court component / PC)	976 / 898	2'515 / 1'242	1'315 / 973
Headline level elasticity vs. tariff (court component / PC)	0.89 / 0.83 ^c	≈0.70 / 0.70	0.79 / 0.88
Party compensation zero (%)	≤8 ^a	12-60 (all five cantons)	27
Appeal success (%)	n.a.	≈29 ^d	18.4 / 17.6 ^d
Lower-instance costs revised on full grant (%)	n.a.	64.0 ^e	65.6 / 52.3 ^e
Ex-ante R^2 (court component / PC)	0.94 / 0.85 ^b	0.63 / 0.60	0.76 / 0.72
Median absolute forecast error (court component / PC)	±21% / ±26%	n.r.	≈ ±23% / ±20%
90% prediction band (factor)	n.r.	≈11	≈4

PC means party compensation. The court component is court costs in Parts I and II and the court fee in Part III. Elasticities use each part's own tariff baseline and specification, so they compare magnitudes only indicatively. n.a. means not applicable and n.r. means not reported. ^a Part I gives an upper bound for cases without a quantified award, not a directly measured zero share. ^b Part I's fit is largely mechanical. ^c Part I treats compensation estimates as indicative. The slope loses significance without copyright cases. ^d Part III gives formal outcome first and loss-share success second. Part II uses the appellant's loss share below one. ^e Part III gives full-grant revision first and loss-share-zero revision second. Part II uses the latter condition. *Source:* Own compilation from Part I, paras. 3 f., 29-35, 74, 131, 142-146, 150, 196 and 198; Part II, paras. 3-7, 15 f., 31-39, 67, 80-84, 89-93 and 117 and Appendix Tables 13, 17 and 29; and this paper.

Table 8: The series at a glance: the three instances compared.

[115] Compression was first documented at the commercial courts in Part I.⁹⁸ It recurs at the cantonal high courts in Part II.⁹⁹ It also appears at the Federal Supreme Court here (Section 5.2). Table 8 aligns the three parts' key figures.¹⁰⁰ Implied elasticities remain between roughly one half and nine tenths even though the tariff regime changes across the three parts.¹⁰¹ Part I uses two heterogeneous cantonal regimes.¹⁰² Part II uses five.¹⁰³ This part uses a single federal schedule. The series supports a recurring direction.¹⁰⁴

[116] In instance order, the court-cost elasticities reported in Parts I and II are 0.89 at first instance and about 0.69 at the cantonal appellate level, while the court-fee elasticity here is 0.79.¹⁰⁵ The middle instance therefore shows the steepest compression. The ladder is not monotone. Three qualifications make this ordering indicative. Headline specifications differ across the parts.¹⁰⁶ Part II's headline excludes realised-effort controls, which lower the Part II elasticities to 0.56 and 0.63 when added, whereas Parts I and III already report fully controlled models.¹⁰⁷ The

98 KÖNIG, Part I (Fn. 5), paras. 6 and 150.

99 KÖNIG, Part II (Fn. 5), paras. 5 f. and 67.

100 The Part I entries are reported in KÖNIG, Part I (Fn. 5), paras. 3 f., 29-35, 74, 131, 142-146, 150, 196 and 198. The Part II entries are reported in KÖNIG, Part II (Fn. 5), paras. 3-7, 15 f., 31-39, 67, 80-84, 89-93 and 117; Appendix Tables 13, 17 and 29.

101 KÖNIG, Part I (Fn. 5), para. 150; KÖNIG, Part II (Fn. 5), para. 67; Section 5.2.

102 KÖNIG, Part I (Fn. 5), para. 74.

103 KÖNIG, Part II (Fn. 5), paras. 15 f.

104 KÖNIG, Part I (Fn. 5), paras. 74 and 150; König, Part II (Fn. 5), paras. 15 f. and 67; Section 5.2.

105 KÖNIG, Part I (Fn. 5), para. 150; KÖNIG, Part II (Fn. 5), para. 67; Section 5.2.

106 KÖNIG, Part I (Fn. 5), para. 150; KÖNIG, Part II (Fn. 5), para. 67; Section 5.2.

107 KÖNIG, Part II (Fn. 5), para. 67; on the Part III specification, see Section 5.2.

samples range from published cantonal judgments to a federal quasi-census.¹⁰⁸ Part I treats its party-compensation estimates as no more than indicative.¹⁰⁹

[117] The series finding therefore rests primarily on the court-side component.¹¹⁰ The same less-than-proportional, effort-graded pattern survives changes in tariff regime, court and instance.¹¹¹ A cantonal-tariff explanation cannot account for the federal result because the schedule here is a single federal one. It did not change throughout the observation window (Section 4.1). Recurrence across regimes is therefore stronger evidence of a general feature of judicial cost-setting than of a trait of one schedule.¹¹²

[118] The exclusion analysis adds one caution to the cross-instance comparison. Removing appeals that adjudicated a purely procedural question moved the court-fee estimate from 0.72 to 0.80. The party-compensation estimate moved from 0.84 to 0.89. Adding first-instance proceeding type to the controls settled them at 0.79 and 0.88. The single exclusion accounts for 86.7 per cent of the court-fee movement and 92.0 per cent of the party-compensation movement. The two smaller exclusions account for the remainder (Section 3.1).

[119] The removed appeals combine a high value in dispute with slight adjudicative effort. Leaving them in exaggerates measured compression. Parts I and II did not draw this distinction.¹¹³ Their estimates may be affected by the same confounding. In that event, they may overstate the magnitude of compression. This possibility cannot be settled here. Doing so would require the same full-text adjudication of their corpora.

[120] The reported magnitude concerns the typical award in the log-scale sense of Section 4.2. It also concerns the conditional median (Section 5.7). A zero-inclusive Poisson estimator instead targets the conditional mean. With the same controls, its level elasticity is 0.83 for the court fee and 0.82 for party compensation (Section 5.7). Both are below one. Compression is therefore not confined to the typical award. It also holds for the conditional mean in levels when the two estimators condition on the same covariates.

108 KÖNIG, Part I (Fn. 5), paras. 29-33; KÖNIG, Part II (Fn. 5), paras. 31-39; Section 3.

109 KÖNIG, Part I (Fn. 5), paras. 131 and 196.

110 KÖNIG, Part I (Fn. 5), para. 150; KÖNIG, Part II (Fn. 5), para. 67; Section 5.2.

111 KÖNIG, Part I (Fn. 5), paras. 74, 150 and 195; KÖNIG, Part II (Fn. 5), paras. 15 f., 67, 78 and 116; Sections 5.2 and 5.3.

112 The mechanism remains untested. Part I documents both compression and a strong association between procedural complexity and court-cost deviations (KÖNIG, Part I [Fn. 5], paras. 150 and 195). Part II documents recurrence under five further cantonal regimes but retains an observational interpretation (KÖNIG, Part II [Fn. 5], paras. 15 f., 67 and 116). The federal schedule likewise directs the Court to the case's scope and difficulty alongside the value in dispute (Section 2.2). None of these results establishes how effort changes with the stakes.

113 Their reported sample and regression designs contain no corresponding full-text exclusion: KÖNIG, Part I [Fn. 5], paras. 34 f. and 72; KÖNIG, Part II [Fn. 5], paras. 34 f. and 59.

	Court-fee compression slope	Party comp. with/ without response (%)	Appeal success (%)
<i>A. By language</i>			
German	0.82	87 / 17	19
French	0.79	93 / 14	19
Italian	0.83	96 / 15	11
<i>B. By civil-law division</i>			
First Civil Law Division (4A)	0.85	92 / 25	19
Second Civil Law Division (5A)	0.66	67 / 1	18

Compression slope: implied court-fee level elasticity ($1 + \hat{\beta}$), positive awards. *Source:* Own calculations.

Table 9: Convergence of cost-setting across languages and divisions at the Federal Supreme Court.

[121] A single court applies the federal tariff but does so in three languages and two divisions. The direction of court-fee compression is uniform across them. The case-characteristics specification is estimated within each subgroup. Language elasticities are about 0.82 in German, 0.79 in French, and 0.83 in Italian. The First Civil Law Division estimate is 0.849 and the Second Civil Law Division estimate 0.663. Within the cost-deviation sample, the First Civil Law Division has 1'138 cases and a median value of CHF 127'045. The Second Civil Law Division has 201 cases and a median of CHF 154'271. The difference is not explained by lower stakes in the Second Civil Law Division's sample.

[122] Within the component-specific court-fee samples, estimates are below one in each language and both divisions. Their common direction is descriptive. The subgroup table reports neither equality tests nor estimates for appeals without an available value.¹¹⁴ Participation and outcome margins also recur across the federal docket. In every language and division, positive party compensation is far more common where a response was invited than where it was not (Section 5.4). Appeal success is roughly flat across the divisions and lower only in Italian (Section 5.5).

[123] Magnitudes vary across languages and divisions. The pattern does not (Table 9). The federal tariff remained unchanged across the fifteen years (Section 4.1). No tariff revision therefore confounds the time dimension. Year fixed effects absorb common year-specific level shifts. A proceeding takes a median of about six months, or 174 days. Non-entry decisions are shortest at 50 days. Successful appeals are longest, ranging from 226 days for remand to 250 days for partial grant.

[124] Within the band, observed position is associated with procedural-effort indicators that mirror the tariff criteria. The exchange of submissions is the strongest margin (Section 5.3). The single-judge inadmissibility track (Section 2.3) sits below the merits docket. Such decisions remain outside the deviation sample as a distinct non-merits procedural path. Their party compensation may be reduced, while the general court-fee criteria do not mechanically place every award near the band floor.

[125] Party compensation is rarely awarded where no response was invited, in about a sixth of appeals (Section 5.4). This is because the respondent rarely incurs compensable effort. Exceptions occur in both directions (Section 5.4). Where a public entity prevails in its official sphere, compensation is withheld as a rule (Section 2.2). These margins are mostly corner solutions. Party compensation is therefore modelled as a hurdle.

¹¹⁴ Party compensation, awarded far less often, is too sparse in the smaller subgroups to weigh on the same basis. Its Italian slope turns mildly positive on only 109 awards.

[126] Two findings have no counterpart at the first instance of Part I and both concern outcome.¹¹⁵ The appeal succeeds only rarely, at 18.4% on the formal outcome measure (Section 5.5). On Part II's measure of the appellant's loss share below one (Section 5.5), the comparable Part III rate is 17.6%.¹¹⁶ This remains well below the cantonal appellate rate of about 29%.¹¹⁷ The Swiss Federal Supreme Court affirms the contested decision far more often.

[127] That pattern is consistent with cognition confined to law and bound to the facts found below (Section 2.1). On the comparable loss-share measure, the success rate in Part III ranges by legal subject from 25.3% in insurance to 12.8% in property (Section 5.5). Part II's range is wider, from 53% in family to 13% in tenancy.¹¹⁸ The second finding concerns lower-instance court costs and party compensation (Section 2.2).

[128] Part II identified the same two-component cost level one instance below.¹¹⁹ At both levels, revision of that cost level is closely associated with appeal success.¹²⁰ Conditional on the appellant's loss share being zero, revision occurs in 52.3% of Part III cases and 64.0% of Part II cases.¹²¹ This comparison uses the same loss-share metric in both parts.¹²² Conditional instead on the formal category of full grant, the Part III rate is 65.6% (Section 5.6).

[129] What the apex changes is incidence. Because it affirms more often, the triggering outcome is rarer and the overall revision rate falls to 10.7% (Section 5.6), below the canton-specific rates reported in Part II.¹²³ Conditional on zero loss share, revision still follows the merits in the same direction at both appellate levels, although the rate is lower here.¹²⁴ This comparison does not establish an identical mechanism across the studies.¹²⁵

[130] The published judgments studied in Parts I and II may not represent the underlying caseload.¹²⁶ This publication-selection limitation does not arise in the same form here. Because the Federal Supreme Court publishes essentially all final decisions, the corpus captures 97.8% of civil appeals disposed of according to the Court's own statistics (Section 3). Unlike the earlier parts, these results are therefore not conditioned by an official selection of decisions for publication.¹²⁷

[131] The deviation estimates nonetheless use only merits appeals with an available value in dispute, about 30% of the merits docket. This value-availability selection need not yield a random sample. The analysis is therefore descriptive and associative. The low share reflects drafting practice.

115 KÖNIG, Part I (Fn. 5), paras. 3 f.

116 KÖNIG, Part II (Fn. 5), para. 84.

117 KÖNIG, Part II (Fn. 5), para. 84.

118 KÖNIG, Part II (Fn. 5), para. 86.

119 KÖNIG, Part II (Fn. 5), para. 65.

120 KÖNIG, Part II (Fn. 5), paras. 82 and 117; Section 5.6.

121 KÖNIG, Part II (Fn. 5), para. 117.

122 KÖNIG, Part II (Fn. 5), para. 117; Section 5.6.

123 KÖNIG, Part II (Fn. 5), para. 82 and Appendix Table 17.

124 KÖNIG, Part II (Fn. 5), paras. 82 and 117; Section 5.6.

125 KÖNIG, Part II (Fn. 5), paras. 82 and 117; Section 5.6.

126 KÖNIG, Part I (Fn. 5), paras. 32 and 190; KÖNIG, Part II (Fn. 5), paras. 31 and 108 f.

127 KÖNIG, Part I (Fn. 5), paras. 32 and 190; KÖNIG, Part II (Fn. 5), paras. 31 and 108-109; Section 3.

[132] For appeals against final decisions, the value-in-dispute requirement for admissibility requires only that the applicable statutory threshold be met (Section 2.1). The admissibility value is determined by the relief sought that remained at issue before the lower instance (Section 2.1). The admissibility consideration therefore typically records that the threshold is exceeded without quantifying the value.¹²⁸ The extraction across the pre-selection frame permits this selection to be examined among its 6'575 pecuniary merits judgments (Appendix Table 62). Selection varies by matter type.

[133] Judgments with an available value in dispute are disproportionately contract and employment claims in the First Civil Law Division. Other judgments more often concern family and tenancy matters in the Second Civil Law Division, but court-awarded amounts are extracted regardless of whether a value is available. They are close across the two groups. Among positive awards, the median court fee is CHF 3'500 in judgments with an available value and CHF 4'000 in judgments without one. The corresponding party-compensation medians are CHF 4'630 and CHF 5'107. Selection thus operates on value availability and matter type.

[134] Among positive awards, judgments without an available value have medians close to and modestly above those in the broader group with an available value in dispute. This comparison partially bounds the concern at one point of the positive-award distributions. It does not test zero shares, tails, dispersion, or conditional equality. It therefore does not establish that the complete award level is unselected. The elasticity is read as conditional on an available value in dispute. Whether the award-to-value relation differs where no value is available remains unidentified from these data.

[135] A validated but not error-free LLM-assisted extraction recovers the court-awarded amounts and value in dispute. The value in dispute is the least reliable regression input and its gaps predominantly reflect omission (Section 3.2). Two design features further bound the reading (Section 4). Because classical measurement error in the tariff anchor would attenuate award elasticity towards zero, the regression is read alongside model-free medians. Those summaries corroborate the descriptive direction. They do not purge error from the value or tariff anchor. In addition, only about ten presiding judges remain, so the wild-cluster bootstrap carries the inference (Section 5.7).

[136] The decision type creates a further limitation. For preliminary and interim decisions, the baseline may follow the pending main matter's value, though only an incidental issue and its associated effort are before the Court.¹²⁹ The estimated deviation may therefore combine within-band cost-setting with a mismatch between statutory value and federal work. The analysis does not identify these channels separately. Even where the analysis can reconstruct a tariff anchor, the judgment does not necessarily show that the Court knew or relied on that amount. A midpoint deviation therefore does not itself identify a discount from a consciously applied value anchor.

¹²⁸ A manual audit of a random sample of the merits judgments without an extracted value confirms this. Most audited judgments state only that the threshold is met or do not quantify the value. Figures stated explicitly but missed by the extraction occur in about three per cent of the audited judgments. That audit counts only figures written in the judgment. The codebook uses a wider construct and also assembles the appeal value from the claims still in dispute. Omissions are therefore more frequent on that construct. Held-out validation puts field recall at about one half (Section 3.2). The two figures measure different things and are not in conflict.

¹²⁹ BGer, 5A_664/2010, 1 June 2011, consid. 7.

8. Conclusion

[137] This paper brings the three-part inquiry into court-awarded litigation costs to the Federal Supreme Court.¹³⁰ It draws on a near-complete census of civil appeals decided by the Federal Supreme Court between 2011 and 2025 in German, French and Italian under a single federal cost regime (Section 3). Within that setting, the central result is clear. Among merits appeals with an available value in dispute, the typical awards exhibit tariff compression (Section 5.2). They rise less than proportionally with their tariff anchors.

[138] Compression operates at two levels. The tariff anchors themselves rise less than proportionally with the value in dispute in their respective component-specific estimation samples (Section 2.2.3). The typical award then rises less than proportionally with that already-degressive anchor. The implied level elasticities are about 0.79 for the court fee and 0.88 for party compensation (Section 5.2). Measured directly against the value in dispute, the award elasticities are about 0.42 and 0.43 (Section 5.2). Both headline estimates are significant under a wild-cluster bootstrap by presiding judge and remain stable across the robustness checks (Sections 5.2 and 5.7). For both cost components, the term typical award denotes the conditional geometric mean (Section 4.2). A zero-inclusive Poisson model with the same controls likewise returns conditional-mean elasticities below one (Section 5.7).

[139] Compression does not exhaust the findings. Position within the tariff band is associated with observable procedural effort, as in both predecessor studies.¹³¹ In the federal analysis, the invitation to respond supplies the sharpest margin (Section 5.3). Party compensation is frequently not awarded at all and this corner solution is closely associated with whether the opposing party was invited to respond (Section 5.4). Revision of lower-instance costs likewise follows the outcome on the merits and occurs almost exclusively when the appeal succeeds (Section 5.6).

[140] Set beside its predecessors, the federal result completes a recurring pattern across all three instances.¹³² The pattern spans heterogeneous cantonal tariffs and the single federal regime. It appears in the court-cost measures reported in Parts I and II and in the federal court fee analysed here (Table 8).¹³³ For party compensation, Part I's estimate remains indicative.¹³⁴ The evidential base thus broadens from published cantonal subsets to a federal quasi-census.¹³⁵

[141] The apex also changes what can be observed. Its near-complete corpus largely removes the official publication choice that constrains the cantonal studies, although the compression estimates remain conditional on component-specific positive-award samples drawn from merits appeals with an available value in dispute (Sections 3 and 3.1).¹³⁶ The Federal Supreme Court's near-complete publication is in this respect exemplary and shows what comprehensive judicial

130 On the conclusions of the predecessor studies, see KÖNIG, Part I (Fn. 5), paras. 194-198 and KÖNIG, Part II (Fn. 5), paras. 115-120.

131 KÖNIG, Part I (Fn. 5), paras. 7 and 123; KÖNIG, Part II (Fn. 5), paras. 7 and 78.

132 See KÖNIG, Part I (Fn. 5), para. 150 and KÖNIG, Part II (Fn. 5), para. 67.

133 KÖNIG, Part I (Fn. 5), paras. 74 and 150; KÖNIG, Part II (Fn. 5), paras. 15 f. and 67; Section 2.2 above.

134 KÖNIG, Part I (Fn. 5), paras. 131 and 196.

135 KÖNIG, Part I (Fn. 5), paras. 150 and 190; KÖNIG, Part II (Fn. 5), paras. 67 and 108 f.; Section 3.

136 KÖNIG, Part I (Fn. 5), paras. 32 and 190; KÖNIG, Part II (Fn. 5), paras. 31 and 108 f.

transparency can offer empirical research. Cantonal courts, whose publication remains uneven and often incomplete, could reach at least a largely comparable standard.¹³⁷

[142] At each instance, the tariff and the observed case profile provide a useful but bounded ex-ante anchor for a positive award.¹³⁸ At the Federal Supreme Court, the in-sample median absolute error is about one quarter in the applicable component-specific model sample (Section 6). For the court fee, the forecast covers only fees assessed by value. For both components, it predicts the amount of a positive award and not the separate probability of a zero award.

[143] One apparent recurrence requires closer parsing. Procedural participation is associated with positive party compensation at the response stage in both appellate studies, but they observe different events.¹³⁹ Part II records whether the respondent filed an appeal response. Positive compensation occurs in 80% of cases with a filed response and in 8% without one.¹⁴⁰ Part III instead records whether the Federal Supreme Court invited a response (Section 2.3) and cannot reliably reconstruct whether a response was ultimately filed. Positive compensation occurs in 90% of cases with an invitation and in 16% without one (Section 5.4). The two results therefore point in the same direction of procedural selection without identifying the same mechanism. Revision of lower-instance costs follows the outcome at both appellate levels.¹⁴¹

[144] Nor does the breadth of the federal corpus remove the series' common limits. All three parts cover adjudicated cases only. Disputes that settle or are withdrawn before judgment therefore fall outside their records.¹⁴² Cases never appealed remain outside Parts II and III but may appear in Part I if they were adjudicated at first instance.¹⁴³ The trilogy therefore concerns court-awarded litigation costs in adjudicated cases.¹⁴⁴ All three analyses are descriptive, estimate associations conditional on their observed records and do not identify causal tariff effects.¹⁴⁵

[145] Four extensions would carry the inquiry beyond the trilogy.¹⁴⁶ More complete and comparable cantonal publication would widen the appellate base beyond Zurich.¹⁴⁷ Tariff revisions would supply variation unavailable under the unchanged federal schedules (Section 4.1). Zug's revised appellate assessment base, in force since 2026, provides one example.¹⁴⁸ Linking the instances at case level would follow a single dispute from the first judgment through appeal to the Federal Supreme Court and would replace three cross-sectional views with that dispute's cost

137 The author is aware that, because the scope of review narrows as a case moves up through the instances, a more reliable publication output at the lower levels would at times involve a far from negligible redaction burden. The Zurich courts in particular nevertheless show that at least an active rate of publication is not impossible at the lower instances. It should equally be acknowledged that there may be a causal relationship between the population size of the Canton of Zurich and its caseload, so that this benchmark is only relative when applied to less populous cantons.

138 KÖNIG, Part I (Fn. 5), paras. 143-147 and 198; KÖNIG, Part II (Fn. 5), paras. 89-93 and 107.

139 For Part II, see KÖNIG, Part II (Fn. 5), paras. 56 and 81; for Part III, see Section 5.4.

140 KÖNIG, Part II (Fn. 5), paras. 56 and 81.

141 KÖNIG, Part II (Fn. 5), paras. 82 and 117; Section 5.6.

142 KÖNIG, Part I (Fn. 5), paras. 4 and 34 f. KÖNIG, Part II (Fn. 5), paras. 4 and 34 f. Section 3.1.

143 On the respective instance and sample definitions, see KÖNIG, Part I (Fn. 5), paras. 3 f. KÖNIG, Part II (Fn. 5), paras. 3 f. and Section 3.1.

144 KÖNIG, Part I (Fn. 5), paras. 3 f. and 34 f. KÖNIG, Part II (Fn. 5), paras. 3 f. and 34 f. Section 3.1.

145 KÖNIG, Part I (Fn. 5), para. 182; KÖNIG, Part II (Fn. 5), para. 116; Sections 3.1 and 4.2.

146 For the predecessor studies' own extensions and limitations, see KÖNIG, Part I (Fn. 5), paras. 182-193 and KÖNIG, Part II (Fn. 5), paras. 114-120.

147 On the published appellate base, see KÖNIG, Part II (Fn. 5), paras. 31-39 and 108 f.

148 KÖNIG, Part II (Fn. 5), para. 21.

trajectory.¹⁴⁹ An international comparison of rigid statutory fee schedules and discretionary cost regimes could test whether court-awarded litigation costs in other jurisdictions also compress against their benchmarks.

[146] The governing criteria tie the award not only to the stakes but also to the effort a case demands (Section 2.2). That architecture is consistent with the pattern observed within the Swiss tariff tradition (Sections 5.2 and 5.3). The trilogy cannot establish whether compression is a general feature of judicial cost-setting or a trait of that tradition.¹⁵⁰ What it can establish is more practical. At every level of the Swiss judicial pyramid, the tariff sets the scale, observable effort helps locate the award within it and a substantial share of the remaining variation lies beyond the models.¹⁵¹ For litigants, the tariff therefore serves at least as an anchor.

149 On the instance and sample definitions of the three cross-sections, see KÖNIG, Part I (Fn. 5), paras. 3 f. KÖNIG, Part II (Fn. 5), paras. 3-7; and Sections 3 and 3.1.

150 On the corresponding limits of causal interpretation, see KÖNIG, Part I (Fn. 5), para. 182 and KÖNIG, Part II (Fn. 5), para. 116.

151 KÖNIG, Part I (Fn. 5), paras. 9, 150 and 195; KÖNIG, Part II (Fn. 5), paras. 67, 78 f. and 116; Sections 5.2 and 5.3.

Appendices

Appendix 1: Coded Variables

The 63 coded variables are grouped by extraction target. The following table lists the extracted variables. The analysis panel also carries computed columns, produced during the run. They include the tariff baselines, decision year, legal-subject grouping, and population metadata. Six further columns record an author decision. They are listed separately in Table 11 below. Reasoning fields used internally during extraction are not stored and are omitted. In the Type column, *id* = identifier, *code* = coded category, *num* = numeric, *int* = integer, *bool* = boolean, *text* = free text, *date* = date.

Variable	Type	Description
A. Identification (pre-filled from the corpus draw)		
<code>case_id</code>	id	Docket number (e.g. 4A_123/2020).
<code>division</code>	id	Civil Law Division (First or Second).
<code>prefix</code>	id	Docket prefix (4A or 5A).
<code>decision_date</code>	id	Date of the judgment.
<code>language</code>	id	Language of the judgment (de/fr/it).
B. Decision and disposition		
<code>decision_type</code>	code	Urteil / Beschluss / Verfügung. Combined ('Urteil und Verfügung') → Urteil.
<code>disposition_type</code>	code	Sachurteil (merits) / Nichteintreten (non-entry) / Abschreibung (struck off: withdrawal, settlement, mootness). The disposition type drives the analysis filter. It also drives the tariff reduction for party compensation (Art. 8(3) FSPCR). The court fee has no automatic reduction. It is instead assessed under Art. 65(2) FSCA within the range set by Art. 65(3)(b) FSCA and item 1 FSFT.
C. Instance chain		
<code>lower_instance_court</code>	text	The court whose decision is contested, verbatim. It may be a cantonal appellate court, a cantonal court acting as sole instance, or the Federal Patent Court.
<code>first_instance_court</code>	text	The court of first instance, if named. In a sole-instance case it may coincide with the court whose decision is contested.
<code>first_instance_date</code>	date	Date of the first-instance decision, if named.
<code>first_instance_proceeding_type</code>	code	CONTROL variable (not a filter): ordentlich / vereinfacht / summarisch (Art. 248 et seq. CPC, including debt-enforcement, register, exequatur and non-contentious matters) / besonderes (special family proceedings, including divorce under Art. 271 et seq. CPC) / nicht_erkennbar.
<code>appeal_outcome</code>	code	Gutheissung / teilweise_Gutheissung / Abweisung / Nichteintreten / Rückweisung. For an allowed appeal, Art. 107(2) FSCA provides for a decision on the merits or remittal. Art. 108(1)(a)-(c) FSCA lists three grounds for non-entry. Dismissal has no separate statutory anchor.
<code>joined_proceedings_indicator</code>	bool	True if joined with a parallel proceeding, or the opposing party filed its own appeal against the same lower-instance decision (substitute for the missing cross-appeal).
D. Values in dispute		
<code>first_instance_dispute_value</code>	num	Value in dispute before the first instance, nominal CHF, if stated.

continued on next page

Table 10 (continued)

Variable	Type	Description
appeal_dispute_value	num	Amount still in dispute before the Federal Supreme Court, reconstructed from the claims the judgment recites. Central tariff driver and baseline anchor. Foreign currency: take the CHF amount the court used. Otherwise note it in dispute_value_reasoning.
dispute_value_reasoning	text	Stored column: short basis or reference for the value in dispute, including any foreign-currency note.
E. Reviewed lower-instance costs		
first_instance_court_costs_nominal	num	Lower-instance court costs as stated in the decision, nominal CHF.
first_instance_party_compensation_nominal	num	Lower-instance party compensation, nominal CHF.
lower_costs_modified_indicator	bool	True if the Federal Supreme Court reallocates the lower-instance costs (Art. 67 and Art. 68(5) FSCA).
F. Federal costs (primary object)		
court_fee_nominal	num	Court fee (Art. 65(2) FSCA) from the operative part, nominal CHF. It is empty on a remittal (Rückweisung) with reserved costs, in which case problem_indicator is also set.
court_costs_total_nominal	num	Total court costs including disbursements (Art. 65(1) FSCA). This is usually equal to court_fee_nominal.
party_compensation_nominal	num	Party compensation (Art. 68 FSCA) from the operative part, nominal CHF, as awarded (including VAT). The value is zero if no compensation is awarded or if the awards are set off. This is not the state compensation under Art. 64 FSCA.
vat_included_indicator	bool	True if the party compensation includes VAT (foreign-domiciled recipient → usually no VAT, cf. domicile_CH).
petitioner_losing_share	num	Cost share borne by the appellant, 0.0-1.0 (Art. 66(1) FSCA). Full dismissal → 1.0.
respondent_losing_share	num	= 1 - petitioner_losing_share.
G. Panel and judges		
panel_size	int	Number of participating judges: 1 (Einzelrichter, Art. 108 FSCA), 3 (rule, Art. 20(1) FSCA), or 5 (Art. 20(2) FSCA). Count the judges in the Besetzung, not the clerk.
fundamental_legal_question_indicator	bool	True only if the court affirms a Rechtsfrage von grundsätzlicher Bedeutung (Art. 74(2)(a) FSCA). Not a mere party assertion or a rejection.
procedure_type	code	Track: ordentlich (circulation) / einzelrichter (Art. 108 FSCA) / dreier_summarisch (Art. 109 FSCA). Not the deliberation form (that is public_hearing_indicator).
presiding_RAW	text	Presiding judge (Abteilungspräsident). Surname, lowercase, spaces→'_'.
judge_no1_RAW	text	Participating judge 1. Same format.
judge_no2_RAW	text	Second sitting federal judge, raw surname.
judge_no3_RAW	text	Third sitting federal judge, raw surname.
judge_no4_RAW	text	Fourth sitting federal judge, raw surname.
judge_no5_RAW	text	Fifth sitting federal judge, raw surname.
clerk_RAW	text	Gerichtsschreiber (fr: Greffier, it: Cancelliere). Surname, lowercase.
H. Procedural-effort indicators		

continued on next page

Table 10 (continued)

Variable	Type	Description
response_invited_indicator	bool	True if a response to the appeal was invited (Art. 102(1) FSCA).
second_exchange_indicator	bool	True if a second exchange of written submissions took place (Art. 102(3) FSCA).
additional_filings_indicator	bool	True if unsolicited further submissions were filed (right to reply). This is distinct from an ordered second exchange.
public_hearing_indicator	bool	True if a party hearing (Art. 57 FSCA) or a public oral deliberation (Art. 58(1) FSCA) took place instead of circulation of the file.
suspensive_effect_indicator	bool	True if the instructing judge actually decided a request for suspensive effect or a provisional measure, granted or dismissed (Art. 103(3) and Art. 104 FSCA). The ex-lege effect under Art. 103(2)(a) FSCA does not count.
instruction_evidence_indicator	bool	True if the court took evidence of its own (Art. 55 FSCA). This is rare at this level because of the binding findings of fact and the bar on new evidence.
duration_days	int	Length of the federal proceeding in days (decision_date minus appeal filing date). It enters the main specification from M4 as a realised-effort covariate.
decision_pages	int	Legacy schema field, never populated (0 of 9'426 rows). The decision length used in the models is text_len.
I. Parties		
petitioners_count	int	Number of appellants (excl. attorneys).
petitioners_attorneys_count	int	Number of appellant-side attorneys (titled lawyers only).
respondents_count	int	Number of respondents (excl. attorneys).
respondents_attorneys_count	int	Number of attorneys acting for the respondents.
petitioner_no1_name_text	text	Plain-text name of the principal appellant. Coded in 188 of 9'426 rows (1.99 per cent) as an identification aid. No analysis uses it. Held internally for data-protection reasons.
petitioner_no1_legal_form	code	Rechtsform code 1-33 (Anhang 1).
petitioner_no1_domicile_CH	bool	True if domiciled in CH. It is empty if not determinable from the anonymised text.
petitioner_no1_NOGA1	code	NOGA L1 code 1-23 (Anhang 2).
respondent_no1_name_text	text	Plain-text name of the principal respondent. Coded in 1'104 of 9'426 rows (11.71 per cent) as an identification aid. No analysis uses it. Held internally for data-protection reasons.
respondent_no1_legal_form	code	Legal form of the principal respondent.
respondent_no1_domicile_CH	bool	True if the principal respondent is domiciled in Switzerland. It also carries the VAT normalisation.
respondent_no1_NOGA1	code	NOGA section (first level) of the principal respondent.
J. Legal subject		
legal_subject_RAW	text	Legal subject from the rubrum (de/fr/it), underscores for spaces, e.g. 'Forderung'.
legal_subject_subsequent	text	An approximate finer classification (e.g. contract type) for generic subjects. It is not used for exact matches.
K. Control flags and filters		
pecuniary_indicator	bool	True = pecuniary dispute (regardless of amount). Only true cases enter the regression.
public_entity_party	bool	True if a public entity is a party in its official capacity (Art. 66(4) and Art. 68(3) FSCA: usually no costs or compensation).

continued on next page

Table 10 (continued)

Variable	Type	Description
legal_aid_granted	bool	True if unentgeltliche Rechtspflege was granted (Art. 64 FSCA).
settled_indicator	bool	True = disposed by settlement, withdrawal, or write-off (Vergleich, Rückzug, Abschreibung; Art. 32 and Art. 71 FSCA). Excluded from the regression.
problem_indicator	bool	True = coding problem / borderline case.
problem_comment	text	Short free-text note on the problem, with reference. Filled in 7 of 9'426 rows (0.07 per cent). It is a coding note, not an analysis variable.

Source: Own classification.

Table 10: The 63 coded variables.

The six panel columns below are not extracted. Each column rests on an author adjudication of the judgment text. The columns define the two denominators the paper reports, so they are stated here in full.

Column	True	Meaning
in_scope	1'806	The disposition population: civil appeals the court disposed of by a merits decision or a non-entry decision. Denominator of the disposition analyses.
r44_pure_process	436	The appeal was adjudicated exclusively on a procedural question. It is removed from the cost-deviation sample. The tariff baseline is built from the value in dispute of the main claim, whereas the court sets the fee according to the effort required by the procedural question, so the anchor and the award reflect different quantities. This mismatch deepens the measured compression mechanically. These cases stay in the disposition population, which would otherwise lose 166 of its 176 non-entry decisions.
excl_r_42	54	Excluded from the cost-deviation sample. On review, the first-instance proceeding was classified as other than a full ordinary proceeding.
excl_r_45	10	Excluded from the cost-deviation sample: the matter falls under the child and adult protection category of the proceeding-type pre-filter, which expressly includes adoption. A counter-search over all 1'806 disposition cases confirmed the set (Section 3).
erstinstanzliche_verfuegung	4	Descriptive flag: the appeal reached the court against an administrative order, not by direct action. It is deliberately not a regressor. Four observations would not support an informative separate regressor.
in_sample	1'339	The cost-deviation sample: in scope, decided on the merits, not purely procedural, and not excluded by the two rules above.

Source: Own adjudication.

Table 11: The six adjudicated panel columns.

Appendix 2: Supplementary Result Tables

This appendix collects the result-bearing analysis tables, including results discussed without a main-text table. These cover the nested models, hurdle, and descriptive compression, as well as robustness, outcomes, subjects, costs, and convergence. The cluster-robust diagnostic is deliberately not tabulated. Showing it beside the inverted interval could imply equivalence, and its measured coverage is reported in Section 4. CF means court fee. PC means party compensation. VID means value in dispute. Where a robustness variant starts from the cost-deviation sample, N_CF and N_PC report the respective component-specific estimation counts.

A. Descriptive statistics (Phase 2)

Metric	Value
Cost-deviation sample	1'339
Years	2011-2025
Median value in dispute	129'628
Languages (de/fr/it)	742/466/131
Panel size (all observed sizes)	2: 1, 3: 1157, 5: 181

Source: Own calculations.

Table 12: Cost-deviation sample and selection.

Award	n>0	n=0	median(>0)	p25	p75
Court fee (CF)	1'333	4	5'000	2'500	6'500
Party compensation (PC, VAT-norm.)	973	366	5'555	2'777	7'400

Source: Own calculations.

Table 13: Distributions of court-awarded litigation costs (court fee and party compensation).

VID quartile	Component	n pos	median award-to-baseline ratio
Q1	CF	315	0.67
Q1	PC	228	0.62
Q2	CF	332	1.0
Q2	PC	233	0.56
Q3	CF	335	0.87
Q3	PC	259	0.65
Q4	CF	333	0.74
Q4	PC	253	0.68

Source: Own calculations.

Table 14: Deviation from tariff (median ratios) by value-in-dispute quartile.

Outcome	N	PC determinable	PC>0	PC=0	Share PC=0 (of determinable)
merits	1'630	1'629	1'139	490	0.301
non-entry	176	176	48	128	0.727

Source: Own calculations.

Table 15: Zero party-compensation rates by disposition.

Panel	n pos (CF)	median award-to-baseline ratio (CF)
Two judges	1	1.0
Three judges	1'136	0.83
Five judges	178	0.86

Source: Own calculations.

Table 16: Court-fee compression by panel size.

B. Main regression: full nested ladder (Phase 3)

Model	Component	N	ln baseline	Elasticity (1+b)	panel_5	panel_1	resp_invited	losing_share	R2_adj
M1	CF	1'315	-0.104*** (0.020)	0.896					0.031
M2	CF	1'315	-0.106*** (0.021)	0.894					0.031
M3	CF	1'315	-0.180*** (0.020)	0.82					0.151
M4	CF	1'315	-0.211*** (0.020)	0.789	-0.040 (0.038)		0.074** (0.032)		0.187
M5	CF	1'315	-0.211*** (0.020)	0.789	-0.040 (0.038)		0.074** (0.032)		0.187
M6	CF	1'315	-0.213*** (0.020)	0.787	-0.044 (0.038)		0.082** (0.035)		0.191

Source: Own calculations.

Table 17: Court-fee deviation model, full ladder M1-M6.

Model	Component	N	ln baseline	Elasticity (1+b)	panel_5	panel_1	resp_invited	losing_share	R2_adj
M1	PC	973	-0.029 (0.032)	0.971					0.001
M2	PC	973	-0.031 (0.032)	0.969					0.002
M3	PC	973	-0.126*** (0.029)	0.874					0.075
M4	PC	973	-0.122*** (0.031)	0.878	0.055 (0.054)		0.223* (0.119)		0.111
M5	PC	973	-0.123*** (0.031)	0.877	0.055 (0.054)		0.223* (0.119)	0.003 (0.041)	0.11
M6	PC	973	-0.122*** (0.029)	0.878	0.052 (0.053)		0.240** (0.120)	0.009 (0.040)	0.139

Source: Own calculations.

Table 18: Party-compensation deviation model, full ladder M1-M6.

Variable	Coef (logit)	OddsRatio	p
ln baseline_pe	-0.045	0.956	0.7241
response_invited_indicator	4.31	74.431	<0.0001
joined_proceedings_indicator	-0.364	0.695	0.4347
fip_t vereinfacht	-1.111	0.329	0.0004
fip_t besonderes	-2.124	0.12	0.0110
fip_t nicht_erkennbar	-2.671	0.069	0.0003

The p column is an HC1 asymptotic diagnostic and is descriptive only. Source: Own calculations.

Table 19: Party-compensation hurdle (logit: positive award).

Condition	N	P(PC>0)
Response invited (response_invited=1)	1'026	0.9
No response invited	313	0.16

Source: Own calculations.

Table 20: Positive-award rate by appeal-response invitation.

Component	coef	t_cluster	p_wildcluster	p_webb	n_cluster	g_eff	rademacher_floor
CF	-0.2134	-6.9729	0.0119	0.0050	10	4.3	0.002
PC	-0.1218	-4.7525	0.0115	0.0152	10	3.7	0.002

Source: Own calculations.

Table 21: Wild-cluster bootstrap of the compression coefficient.

C. Robustness checks (Phase 4)

Component	N	PPML elasticity (ln baseline)	SE	p(elast=1)	<1 sig.	Zero share
CF	1'319	0.825	0.024	<0.0001	yes	0.003
PC	1'339	0.818	0.03	<0.0001	yes	0.273

The p column is an HC1 asymptotic diagnostic and is descriptive only. Source: Own calculations.

Table 22: R1: PPML, zero-inclusive conditional-mean elasticity.

Component	Baseline band	n pos	median ln	median award-to-baseline ratio	elast_M6
CF	min	1'315	0.773	2.17	0.862
CF	mid	1'315	-0.182	0.83	0.787
CF	max	1'315	-0.693	0.5	0.762
PC	min	973	0.329	1.39	0.856
PC	mid	973	-0.48	0.62	0.878
PC	max	973	-0.913	0.4	0.87

Source: Own calculations.

Table 23: R2: tariff-band anchor sensitivity (min/mid/max).

Component	N	coef with LS	coef without LS	Delta
CF	1'315	-0.213	-0.208	0.006
PC	973	-0.122	-0.114	0.007

Source: Own calculations.

Table 24: R4: dropping the legal-subject controls.

Variant	N	coef ln baseline (PC)
VAT-normalised (headline)	973	-0.122
nominal (no VAT)	973	-0.121

Source: Own calculations.

Table 25: R5: VAT normalisation on/off.

Variant	N_CF	coef_CF	N_PC	coef_PC
Cost-deviation sample (merits, 30k)	1'315	-0.213	973	-0.122
incl. non-entry (30k, broad)	1'476	-0.303	1'021	-0.16
Threshold 15k (merits; no R-44 adjudication below 30k)	1'440	-0.187	1'108	-0.146

Source: Own calculations.

Table 26: R6: sample and threshold variants.

Variant	N_CF	elast_CF	N_PC	elast_PC
Cost-deviation sample	1'315	0.7866	973	0.8782
less the four administrative-route merits appeals	1'311	0.7881	969	0.8792

These appeals reached the court against an administrative order (author decision 4). Source: Own calculations.

Table 27: R16: headline elasticity without the four administrative-route merits appeals.

Variant	N_CF	elast_CF	adjR2_CF	N_PC	elast_PC	adjR2_PC
M6 as reported, with proceeding type	1'315	0.7866	0.1911	973	0.8782	0.1393
M6 without proceeding type	1'315	0.7953	0.1619	973	0.8873	0.1296

This counter-run implements author decision 5. It prevents an ex-post choice of controls. Source: Own calculations.

Table 28: R17: headline specification with and without first-instance proceeding-type controls.

Variant	set removed	effective removal	N_CF	elast_CF	N_PC	elast_PC
Cost-deviation sample	0	0	1'315	0.7866	973	0.8782
less joined judgments with a single fee	18	17	1'298	0.7869	960	0.8776
less all joined judgments	35	31	1'284	0.7849	951	0.8756

One excludes joint-fee cases. The other excludes all joined judgments. The table distinguishes the provenance set from the number removed from the cost-deviation sample. Source: Own calculations.

Table 29: R18: two counter-runs for judgments joining several appeals.

Variant	N_CF	coef_CF	elast_CF	coef_PC	elast_PC	median_ratio_CF	pct_below_baseline_CF	median_ratio_PC
Cost-deviation sample	1'315	-0.213	0.787	-0.122	0.878	0.833	63.0	0.619
drop value>100m	1'307	-0.227	0.773	-0.103	0.897	0.833	63.4	0.619
drop top value pct	1'301	-0.233	0.767	-0.11	0.89	0.833	63.6	0.619

The row labelled cost-deviation sample reports component-specific estimation counts. *Source:* Own calculations.

Table 30: R7: high-leverage and extreme-value trim.

Component	n_cl	g_eff	dom_pct	coef_full	p_webb	p_rade	coef_drop	p_webb_drop	p_rade_drop	rade_floor
CF	10	4.3	35.5	-0.213	0.005	0.0119	-0.207	0.0173	0.0197	0.002
PC	10	3.7	36.6	-0.122	0.0152	0.0115	-0.1	0.0121	0.0079	0.002

Source: Own calculations.

Table 31: R8: few-cluster robustness check (drop the dominant presiding judge).

Component	N	residual_SD_ln	band90_factor	median_abs_pct_error
CF	1'315	0.384	3.5	22.7
PC	973	0.471	4.7	19.5

Source: Own calculations.

Table 32: R10: ex-ante prediction band.

Component	N	elast_vs_tariff	elast_bl_vs_value	elast_aw_vs_value	elast_diff_aw_bl	identity	same_sample
CF	1'315	0.787	0.48	0.418	-0.062	0.0	yes
PC	973	0.878	0.458	0.43	-0.027	0.0	yes

Source: Own calculations.

Table 33: R11: anchor decomposition, tariff versus value-in-dispute elasticity (arithmetic, no attribution).

Variant	N_CF	N_PC	elast_CF	elast_PC	median_ratio_CF	median_ratio_PC
Cost-deviation sample	1'315	973	0.787	0.878	0.833	0.619
drop value == 30'000	1'261	934	0.771	0.875	0.851	0.619
floor raised to 50'000	1'048	779	0.699	0.907	0.923	0.648

The row labelled cost-deviation sample reports component-specific estimation counts. *Source:* Own calculations.

Table 34: R12: threshold-artefact trim at the admissibility floor.

component	variant	n	dropped	elast	p_webb	elast_lo	elast_hi
CF	component-specific estimation sample	1'315	0	0.7866	0.0048	0.6353	0.8715
CF	b drop not-in-text	1'247	76	0.7884	0.005	0.6604	0.8603
CF	c drop computed sums	1'184	143	0.7899	0.0048	0.6434	0.883
CF	d directly stated only	1'116	219	0.7912	0.0052	0.6717	0.8695
PC	component-specific estimation sample	973	0	0.8782	0.0163	0.7615	0.9426
PC	b drop not-in-text	923	57	0.8773	0.0161	0.7601	0.9349
PC	c drop computed sums	873	111	0.8748	0.0141	0.7329	0.9477
PC	d directly stated only	823	168	0.8738	0.0132	0.7278	0.9396

Variants b and c each start from the respective component-specific estimation sample. Variant d excludes both classes. Webb intervals use 19'999 replications and seed 42. The headline bootstrap instead uses 99'999 replications (Section 4). Test inversion evaluates every grid point with a full bootstrap. The column *dropped* counts provenance cases removed before estimation. The column *n* counts complete cases used for each component. The starting count minus *dropped* need not equal *n*, because some removed cases already lack that cost component. *Source:* Own calculations.

Table 35: R13: separate and combined exclusions of two suspect provenance classes.

component	variant	grid bound	sharpened bound	grid step
CF	component-specific estimation sample	[0.6642, 0.8601]	[0.6353, 0.8715]	0.0122
CF	b drop not-in-text	[0.6783, 0.8545]	[0.6604, 0.8603]	0.011
CF	c drop computed sums	[0.6625, 0.8791]	[0.6434, 0.8830]	0.0127
CF	d directly stated only	[0.6767, 0.8599]	[0.6717, 0.8695]	0.0115
PC	component-specific estimation sample	[0.7756, 0.9397]	[0.7615, 0.9426]	0.0103
PC	b drop not-in-text	[0.7660, 0.9329]	[0.7601, 0.9349]	0.0111
PC	c drop computed sums	[0.7409, 0.9417]	[0.7329, 0.9477]	0.0134
PC	d directly stated only	[0.7409, 0.9328]	[0.7278, 0.9396]	0.0148

Each grid bound is a grid point. Its resolution follows the grid step, not the replication count. *Source:* Own calculations.

Table 36: R13: grid bounds and bisection-sharpened bounds for eight cells.

provenance class	n	flagged	share flagged	separates?
sum_of_two_items	132	24	0.182	no: same rate as directly stated
not_in_text	68	15	0.221	no: same rate as directly stated
stated as a number	1'139	199	0.175	no: same rate as directly stated

Source: Own calculations.

Table 37: R13 diagnosis: the dispute-value provenance flag does not separate the categories.

Component	N	elast_OLS	elast_median_q50	median_OLS_resid
CF	1'315	0.787	0.831	0.026
PC	973	0.878	0.949	0.039

Source: Own calculations.

Table 38: R14: median regression, 0.5 quantile with the full control set.

component	variant	n	dropped	elast	elast_full_precision
CF	drop single-instance	1'238	77	0.7848	0.784831192878218
CF	component-specific estimation sample	1'315	0	0.7866	0.7866065642419101
PC	drop single-instance	943	77	0.8822	0.8821532570604557
PC	component-specific estimation sample	973	0	0.8782	0.8781529735990486

The main text reports the rounded result. Full precision distinguishes it from R13 variant b. The column *dropped* counts cases removed before estimation. The column *n* counts complete cases used for each component. The starting count minus *dropped* need not equal *n*, because some removed cases already lack that cost component. *Source:* Own calculations.

Table 39: R15: headline elasticity without 77 single-instance supplementary-insurance appeals.

D. Ex-ante prediction (Phase 5)

Component	N	R^2 ex-ante M3	R^2 without year FE	R^2 tariff-only M1	R^2 adj. M3
CF	1'315	0.764	0.762	0.706	0.742
PC	973	0.716	0.713	0.654	0.68

Source: Own calculations.

Table 40: Ex-ante R^2 (M3 predictors) and tariff-only R^2 .

VID (CHF)	CF predicted (CHF)	PC predicted, VAT-normalised (CHF)
30'000	2'800	2'440
75'000	2'990	3'946
150'000	4'256	5'750
300'000	5'934	5'750
750'000	9'019	7'955
2'000'000	15'131	10'075

Source: Own calculations.

Table 41: Ready reckoner: central predicted court-awarded litigation costs by value in dispute.

E. Appeal outcomes and legal subjects (Phase 6)

Metric	Numerator	Denominator	Per cent
Appeal success (Appeal outcome)	332	1'804	18.4
Appellant's loss share below one (petitioner_losing_share)	316	1'791	17.6
Observations with both measures		1'791	
Cases where the two measures disagree	32	1'791	1.8
of which joined judgments, verified on the operative part	10	32	
of which flagged joined_proceedings_indicator in the panel	15	32	

The first is outcome-based. The second uses the cost share. *Source:* Own calculations.

Table 42: Two appeal-success measures and their disagreements.

Appeal outcome	N	Share pct
dismissal	1'296	71.8
non-entry	176	9.8
partial	170	9.4
grant	125	6.9
remand	37	2.1

Source: Own calculations.

Table 43: Appeal-outcome distribution.

Division	N	Success rate pct
4A	1'479	18.5
5A	325	17.8

Source: Own calculations.

Table 44: Appeal success rate by civil-law division.

Language	N	Success rate pct
de	994	19.4
fr	619	19.2
it	191	10.5

Source: Own calculations.

Table 45: Appeal success rate by language.

VID quartile	N	Success rate pct
Q1 (low)	452	17.5
Q2	450	17.6
Q3	451	16.6
Q4 (high)	451	22.0

Source: Own calculations.

Table 46: Appeal success rate by value-in-dispute quartile.

Legal subject	N	Success rate pct
insurance	153	28.1
corporate	77	26.0
family	138	22.5
liability	149	21.5
enforcement	43	20.9
employment	208	17.3
inheritance	66	16.7
misc	428	16.6
property	47	14.9
tenancy	67	14.9
contract_claim	419	14.3

Source: Own calculations.

Table 47: Appeal success rate by legal subject.

Legal subject	N	CF0_pct	PC0_pct
misc	430	1.4	34.2
contract_claim	419	1.0	20.0
employment	208	0.5	23.1
insurance	153	2.0	48.4
liability	149	1.3	25.5
family	138	3.6	63.0
corporate	77	1.3	15.6
tenancy	67	1.5	28.4
inheritance	66	3.0	60.6
property	47	0.0	76.6
enforcement	43	0.0	67.4
ip_competition	9	0.0	44.4

Source: Own calculations.

Table 48: Zero court-fee and zero party-compensation shares by subject.

Group	N	CF=0 %	PC=0 %
Public entity is a party (descriptive; cf. Art. 66(4) / 68(3) FSCA)	141	6.4	58.2
other cases	1'665	1.0	32.2

Source: Own calculations.

Table 49: Public-entity involvement, descriptive (cf. Art. 66(4) and Art. 68(3) FSCA).

Sub-matter (keyword cluster)	N	Success pct
contract_claim: sale	49	2
contract_claim: work contract	41	10
family: divorce	20	15
contract_claim: monetary claim	45	16
tenancy: rent / other	38	21
family: maintenance	26	27

Source: Own calculations.

Table 50: Appeal success by fine-grained sub-matter.

F. Extensions: second cost level and out-of-sample ex-ante (Phase 7)

Group	N	Revision rate pct
overall	1'806	10.7

Source: Own calculations.

Table 51: Lower-instance cost-revision rate.

Outcome	N	Revision rate pct
grant	125	65.6
partial	170	60.0
remand	37	24.3
dismissal	1'296	0.0
non-entry	176	0.0

Source: Own calculations.

Table 52: Lower-instance cost-revision rate by appeal outcome.

Success measure	N	Revisions	Revision rate pct
petitioner_losing_share = 0	195	102	52.3

Source: Own calculations.

Table 53: Lower-instance cost-revision rate where the appellant's loss share is zero.

Outcome	no cost revision	cost revision
Unsuccessful (dismissal, non-entry)	1'472	0
Successful (grant, partial, remand)	139	193

The contingency table exhibits quasi-complete separation. No unsuccessful appeal has a cost revision. An unpenalised logit therefore has no finite maximum-likelihood estimate. Fisher's exact test gives $p = 4.8 \times 10^{-169}$ on $n = 1'804$.

Source: Own calculations.

Table 54: Lower-instance cost revision by appeal outcome.

Component	N	R^2 in-sample M3	R^2 5-fold CV M3
CF	1'315	0.764	0.68
PC	973	0.716	0.632

Source: Own calculations.

Table 55: Ex-ante predictability: in-sample versus five-fold cross-validated R^2 .

Component	train_years	test_years	N_train	N_test	R2_out_of_time	median_abs_pct_error	mean_log_error
CF	2011-2021	2022-2025	1'034	281	0.765	25.8	0.05
PC	2011-2021	2022-2025	759	214	0.653	23.1	0.074

Source: Own calculations.

Table 56: Out-of-time ex-ante prediction: training period 2011-2021 and test period 2022-2025.

G. Duration and convergence (Phase 8)

Subset	N	Median (days)	Median (months)
overall	726	174	5.7

Source: Own calculations.

Table 57: Appeal duration, overall.

Division	Median (days)	N
4A	171	413
5A	180	313

Source: Own calculations.

Table 58: Appeal duration by civil-law division.

Outcome	Median (days)	N
non-entry	50	111
dismissal	175	488
remand	226	18
grant	227	47
partial	250	62

Source: Own calculations.

Table 59: Appeal duration by appeal outcome.

Group	Component	coef ln baseline	Elasticity (1+b)	N
All	CF	-0.18	0.82	1'315
All	PC	-0.126	0.874	973
Language=de	CF	-0.183	0.817	725
Language=de	PC	-0.125	0.875	514
Language=fr	CF	-0.211	0.789	459
Language=fr	PC	-0.211	0.789	350
Language=it	CF	-0.172	0.828	131
Language=it	PC	0.115	1.115	109
Division=4A	CF	-0.151	0.849	1'116
Division=4A	PC	-0.097	0.903	914
Division=5A	CF	-0.337	0.663	199
Division=5A	PC	-0.047	0.953	59

Source: Own calculations.

Table 60: Convergence: compression coefficient by language and division.

Group	N (response)	PC>0 with response (%)	N (no response)	PC>0 without response (%)
All	1'026	90.0	313	16.0
Language=de	556	86.7	186	17.2
Language=fr	359	93.3	107	14.0
Language=it	111	95.5	20	15.0
Division=4A	940	92.0	198	24.7
Division=5A	86	67.4	115	0.9

Source: Own calculations.

Table 61: Party-compensation response gateway by language and division.

H. Selection by value-in-dispute availability (Phase 10)

Characteristic	Available value	No available value
Merits judgments (N)	2'007	4'568
Court fee extracted, n (missing)	1'996 (11)	4'539 (29)
Party compensation extracted, n (missing)	2'006 (1)	4'566 (2)
Court fee, median CHF	3'500	4'000
Party compensation, median CHF	4'630	5'107
Response invited, %	72	67
Outcome available, n (missing)	2'004 (3)	4'566 (2)
Appeal success (Appeal outcome), %	20.3	22.1
First / Second Civil Law Division, %	84 / 16	65 / 35
German / French / Italian, %	57 / 34 / 10	46 / 48 / 6
Contract / employment, %	21 / 17	12 / 11
Family / tenancy, %	7 / 5	18 / 11

Cost medians use positive awards. Source: Own calculations.

Table 62: Selection by value-in-dispute availability: merits judgments grouped by whether a value in dispute is available.

ADRIAN KÖNIG, M.Sc., MLaw, CAS ILE, is a research associate and PhD candidate at the Institute for Private International Law and Civil Procedure (CIVPRO) at the University of Bern. His research focuses on data-driven disputes. Inquiries may be addressed via adriankoenig.ch.

Ethics statement. The study analyses court judgments published in anonymised form. Public availability does not eliminate privacy risks, as judgments may contain sensitive personal, financial, family, or health information. The material was used solely for scholarly analysis. The results are reported in aggregate, and no attempt was made to re-identify anonymised individuals. DeepL was used for linguistic revision and translation across all three parts of *From Tariff to Award*. OpenAI's Codex was used for debugging and the verification of calculations across the series. The automated extraction of legal information reported in the papers was performed using OpenAI's GPT-4o, GPT-5.1, and GPT-5.4-mini. The specific applications are described in the respective sections. All three parts of *From Tariff to Award* were written solely by the author, who bears full responsibility for their content. The papers were reviewed by external experts or benefited from their input, all of whom are listed in the Acknowledgements.

Reproducibility statement. The extracted case-level data are not publicly available. The author preserves the fixed source snapshots, codebooks, code, validation records and generated outputs used for the analyses. Researchers seeking to inspect or reproduce a result may contact the author at info@adriankoenig.ch. Access and assistance remain subject to applicable data-protection and source-use constraints.

Acknowledgements. The author thanks Michael Ryf, M.Sc. in Business Administration, PhD candidate in Finance at the Institute for Financial Management, University of Bern; Wjera Yell Leutenegger, M.Sc. in Economics, PhD candidate in Macrofinance at the Department of Economics, University of Bern; Roman Wixinger, M.Sc. in Physics from ETH Zurich, Data Scientist at Ergon Informatik AG in Zurich; Philip Pawlowsky, B.Sc. from ETH Zurich, student of Computational Science and Engineering at ETH Zurich; Jonas Käser, M.Sc. in Economics, working in Financial Risk Management at Entris Banking AG; and Andjelika Mirkovic, stud. iur., for their valuable input and discussions. The author also thanks Jonas Hertner, MLaw, attorney-at-law, for valuable work on the OpenCaseLaw project. The database developed through that project was used in Parts II and III of the series. Further thanks are due to Dr. iur. Meinrad Vetter, LL.M., attorney-at-law and President of the Commercial Court of the Canton of Aargau, for kindly providing data, and to Oliver Graf, MLaw, Media Officer at the High Court of the Canton of Zurich, for assistance with data-related questions concerning the Commercial Court of Zurich. The author also thanks Prof. Dr. Florian Eichel and the Institute for Private International Law and Civil Procedure (CIVPRO) at the University of Bern for their support. This research series originated in the author's master's thesis in corporate finance, completed at the institute directed by Prof. Dr. Philip Valta, whom the author thanks for this formative opportunity.