ARTICLE

A Two-Stage AI-Based Framework for Determining Insurance Broker Commissions in the Healthcare Industry

Jeshwanth Reddy Machireddy[1]

[1] Independent Researcher

Abstract

Insurance brokers play a critical role in connecting consumers and employers with health insurance plans in the United States, and their compensation in the form of commissions significantly influences the dynamics of the healthcare insurance market. This paper proposes a conceptual two-stage artificial intelligence (AI) driven framework for determining insurance broker commissions in the U.S. healthcare industry. The framework is designed to account for multiple variables that are often overlooked by traditional commission structures, including the size of the insurance policy, the risk profile of the client, the characteristics of the healthcare plan, and the historical performance of the broker. In the first stage of the framework, a data-driven model analyzes policy-specific factors to compute a baseline commission recommendation. In the second stage, a subsequent AI model refines this commission by incorporating broker-specific performance metrics, thereby personalizing the compensation to align with the broker's track record and value delivery. This two-stage approach allows for a modular and comprehensive analysis that mirrors real-world commission practices (such as base commissions combined with performance-based bonuses), but it is enhanced through AI to achieve greater precision and adaptability. The paper details the design of each stage and the interaction between them, provides a mathematical representation of the framework, and discusses how the model can handle complex variables inherent to the healthcare insurance commission process. Although presented conceptually without empirical case studies, the proposed framework offers a blueprint for leveraging AI in commission determination, aiming to improve incentive alignment and efficiency in the broker-mediated health insurance marketplace.

Keywords: AI-driven framework, broker compensation, healthcare insurance, performance-based commissions, policy-specific variables, risk profiling, U.S. insurance market

Copyright

Machireddy, J. R. (2022), A Two-Stage AI-Based Framework for Determining Insurance Broker Commissions in the Healthcare Industry. Transactions on Artificial Intelligence, Machine Learning, and Cognitive Systems, 7(6), 1–21.

Submitted: 2022
Accepted: 2022
Published: 2022

Vol. 7, **No.** 6, 2022.

1 Introduction

Insurance broker commissions are a substantial part of the U.S [1]–[3]. healthcare insurance industry's administrative cost and incentive structure [4], [5]. Brokers serve as go-betweens between clients—e.g., employers in search of group health coverage or individuals selecting policies—and carriers, and are normally paid by insurers on a commission basis for sold policies [6]–[8]. These commissions typically take the form of a percentage of the insurance premium or as a flat rate per-insured amount and might be calculated based on the policy category or market niche [9], [10]. For instance, a broker who brokers a small-group health insurance policy will earn a commission of about five percent of premium, whereas in the case of large-group or specialty plans, the commission arrangement can vary or be negotiated on a flat per-member basis [11]–[13]. In whatever

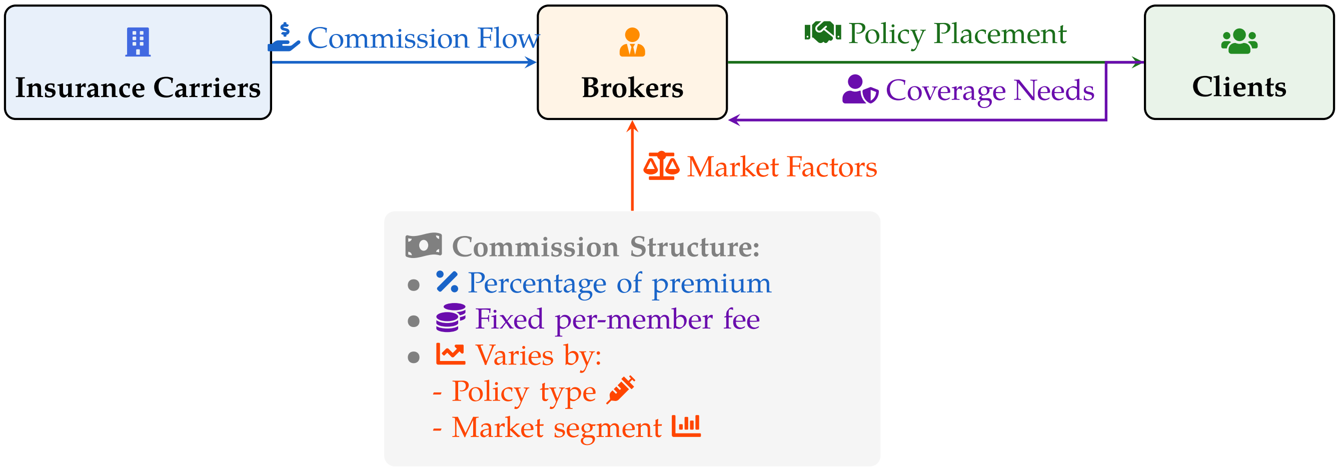

Figure 1. Insurance broker ecosystem showing intermediaries (brokers) connecting carriers and clients, with commission structures based on premium percentages or fixed fees, influenced by policy types and market segments.

particular form, broker commissions are intended to compensate brokers for their work in advising clients, enrolling members, and servicing accounts. [14], [15]

Traditional methods of establishing broker commissions on health insurance lean toward the use of general rules and industry norms [16]–[18]. One-size-fits-all commission grids are typical practice among product-line-based, mainly insurance carrier-based firms (e.g., HMO, PPO, or dental policy) and policy size (policyholder number or total premium) [19], [20]. Such traditional methods, although simple, are not necessarily likely to identify the subtle differences between individual cases [21]–[23]. For instance, two companies might both pay the same amount for a health policy, but one company's employees might have a very different health-risk profile from the other company's employees, maybe requiring more work on plan customization and client relationships for the broker [24], [25]. Under an average commission arrangement, the broker would get the same per-account payment [26], [27]. Thus, an extremely intricate health policy with widespread coverage or a very limited network may require more brokering and explanation from an agent than a simple policy to read, but that intricacy may not correspond to the commission in the way a flat system may not realize [28]–[30].

Shortcomings in existing methodologies for the calculation of commissions indicate a probable market for additional advanced, data-based methods [31]–[33]. Specifically, machine learning and artificial intelligence methods provide a way of adding more variables and historical information to the commission-setting process [34], [35]. By analyzing past insurance sales history, client results, and broker performance data, AI systems can potentially determine patterns that suggest when higher or lower commissions are warranted [36]–[38]. For example, a model AI can learn that policies sold to customers in certain high-risk industries are more likely to require more broker facilitation and thus slightly higher commission will most likely be related to successful plan implementation and retention of the customer in such a case [39], [40]. Likewise, it may observe that very experienced brokers with a good name can handle larger accounts in an efficient manner, which may enable the insurer to remunerate them differently (e.g., via bonus schemes) than less experienced brokers, who may be more dependent on strong financial incentives to produce the same outcomes. [41]–[43]

We suggest a two-stage AI-based model for the calculation of insurance broker commissions that systematically integrates a variety of essential factors into the decision-making process. The model is based on the U.S. health insurance market, acknowledging the industry's nuances such as varied plan types, price restrictions imposed by regulators (e.g., limits on premium variation based on health status under the Affordable Care Act), and the large role brokers have in plan distribution. The first phase of the model determines the specifications of all insurance policies and client cases to generate a preliminary commission suggestion. It uses AI to estimate such factors as policy size (i.e., total premiums or lives covered) risk profile of the client (predictive of future healthcare use or claims risk), and plan features being offered for sale (such as coverage extent, network size, or product novelty). The output of Stage 1 is a baseline commission value which is tailored based on the difficulty and predicted quantity of work involved in

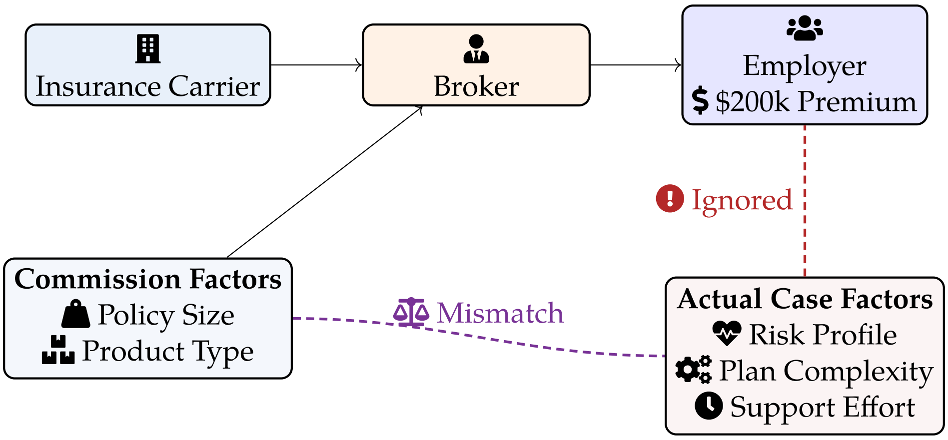

Figure 2. Visualization of commission structure limitations: Standard factors (policy size/product type) fail to account for risk profile, plan complexity, and support effort in individual employer cases.

Commission Formats in Healthcare Insurance

Policy Type	Commission Type	Typical Rate	Risk Level	Broker Effort
Small-Group Plan	Percentage of Premium	5%	Medium	Moderate
Large-Group Plan	Flat Fee per Member	$20/member	Low	Low
Specialized Plan	Negotiated Rate	Varies	High	High
Dental/Vision Add-ons	Fixed Supplement	$10/member	Low	Minimal

Table 1. Common commission formats across healthcare insurance policy types.

the individual case under consideration.

The second phase of the model applies the baseline commission developed in Stage 1 and modifies it based on the previous performance of the individual broker handling the account. This performance-based update also considers aspects such as the broker's historical client retention record, portfolio growth, customer satisfaction, and other metrics of their effectiveness and efficiency. Taking these into account, Stage 2 tailors the commission, i.e., two brokers selling identical insurance policies can have varying commission recommendations if one has a history where they consistently provide more value or needs a different level of incentive. In practice, Stage 2 serves the same purpose as a performance bonus or contingent commission, in addition to the underlying commission: a notion already existing in the majority of brokerage compensation arrangements (e.g., year-end bonuses for meeting sales goals or profit-sharing commissions on low-claim books), but here calculated in a data-driven way by an AI model.

Union of these two periods creates a well-structured framework which puts setting the commission into concordance with both the detail of the insurance transaction as well as the nature of the broker intermediary. This kind of framework is able to render the decision of the commission more sensitive and more individually tailored to the condition of each transaction. Instead of paying every broker the same commission for a given product or by ad-hoc agreement for exceptions, an insurer can employ this technique to create stable but tailored offers of commissions based on anticipated work and worth in every sale. Additionally, through the use of AI, the system is able to keep learning and fine-tuning its suggestions as more information on results (like policy persistency, claim experience, sale conversion ratios, and broker behavior patterns) is published over time.

Throughout the following sections, we begin by offering background on insurance broker commissions in healthcare and outline the key drivers that affect commission structures. Second, we describe the architecture of the two-stage AI framework proposed and provide the justification for the design. We then move on to explain the specifics of every stage of the framework, how the models work, and how they integrate different variables like policy features and broker performance. We then provide a rudimentary mathematical depiction of the framework and demonstrate its application using an example

Factors Influencing Broker Commission Variations

Variable	Impact	Example Scenario	Broker Response
Workforce Risk Profile	Higher Commission	High-risk industry	Intensive effort required
Plan Complexity	Higher Commission	Narrow provider network	Extensive explanation needed
Account Size	Lower Commission	Large employer group	Streamlined handling
Broker Experience	Variable Bonus	Veteran broker efficiency	Adjust compensation strategy

Table 2. Variables influencing broker commissions and recommended broker responses.

sample situation. We then introduce the possible implications and challenges of the proposed approach, and close with concluding remarks, summarizing the contributions of this conceptual framework and possible implications and future research directions for AI-based commission determination research.

2 Background

2.1 Broker Commissions in U.S. Healthcare Insurance

Health insurance brokers play an important middleman role helping employers and individuals navigate a complex web of coverage. Brokers advise clients on selecting a plan, negotiate on behalf of clients with carriers, and help with enrollment and ongoing service issues. Brokers receive compensation for their services primarily from carriers in the form of commissions. These commissions typically are built into the insurance premium and are not paid out of pocket by the client directly (specifically within the employer-sponsored market), but ultimately such costs are incorporated into overall premium rates.

The brokerage commission payment method of healthcare may vary by insurer and market segment. One general model is percentage-of-premium commission. For example, in most small-group and individual health insurance markets, a broker may receive a commission that is a percentage of the monthly premium for each policy sold or renewed. Not only is it commonplace for such premiums to be in the range, for example, of 3% and 7% of the premium for employer-sponsored health insurance, but percentages are established by each insurance company and can be subject to competitive pressures. In others, the percentage may fall as the policy size (and premium) grows, on volume discount grounds – this is true for a very high premium, big group employer with significantly lower commission percentage than for a small group, on the grounds that the work in servicing one large client is less, per-insured, than servicing many small clients to the same volume of premiums. Conversely, very small policies (e.g., a family or individual plan) might pay a higher percentage commission or flat minimum charge so that the broker is adequately remunerated for the fixed work involved in the writing of any policy, however small.

Yet another typical commission arrangement is a flat fee per-member or per-employee, typically expressed as a dollar amount per enrolled member per month. This strategy is sometimes utilized in larger group markets or circumstances where a fee-based compensation is negotiated. An example is when a broker earns a commission of $20 for every employee each month for an employer's health plan regardless of the premium of each employee's selected plan choices. These structures facilitate budgeting for the employer and can decouple the broker's fee from the premium levels, though ultimately insurers still normally finance this out of the premium charges. A few employers and brokers have recently also tried fee-for-service or consulting fee structures (where the employer pays the broker directly for their services in the form of a fee, as opposed to a commission paid by the insurer) and fees based on performance, but the standard remains the traditional commission paid by insurers.

It must be noted that besides the base commissions, insurers usually have additional tiers of broker compensation above. These could be "override" commissions or bonuses based on performance, tied to general measures of performance, such as the volume of business a broker places with the insurer within a year, the growth in that volume, or the broker's book of business loss ratio (claims incurred vs. premiums earned). For instance, an insurer might pay a 1% bonus commission on all sales if the broker insures over a certain number of clients or achieves a high rate of renewal on their portfolio. These incentives are intended to reward brokers for performance and loyalty and to encourage them to bring in more business to the insurer or retain high-value accounts. While these payments are not usually disclosed to customers, they account significantly for a broker's overall compensation package and still underscore that calculating commissions can be a multifaceted exercise with various inputs and factors.

2.2 Factors Influencing Commission Decisions

Insurance companies traditionally consider a handful of significant considerations in setting commission levels, although often on a qualitative or rule-of-thumb basis. Policy size or premium volume associated with the sale is one of the significant considerations. As explained, more premium policies (i.e., covering many people or offering very extensive benefits) generate more commission dollars in absolute terms if a percentage rate is applied, which can justify the application of a lower percentage rate for very large cases to keep the broker's total compensation at reasonable levels. Conversely, small policies might need a higher percentage or a flat minimum fee to allow the broker's compensation to compensate the broker for time and effort. Group or policy size is also sometimes connected with the complexity of service: a large employer can either have multiple locations, huge enrollment support needs, and year-round servicing, or, alternatively, a large employer can have an internal HR department that takes on many tasks, making the broker's job easier in some respects. These nuances are not generally explicitly accounted for in traditional commission arrangements but are implicitly acknowledged through experience and ad hoc modification or exceptions. Under an AI-driven approach, these nuances can be quantifiably captured by examining data on past cases of different sizes and service requirements. A second important factor is the risk profile of the client or expected health spend level. In health insurance, especially for employer group policies, underwriters assess the risk profile of the group (e.g., based on demographic information, industry type, prior claims experience if any, etc.) in order to rate. Even though commissions are not typically adjusted case-by-case for risk in current practice (since they are typically set as a flat percentage or fee), the economics of a low-risk client and a high-risk client are quite different. A high-risk group may be able to command a high premium (to cover expected claims), which would increase commission if it is a flat percentage, but meanwhile the profit margin of the insurer on that policy may be lower due to those higher claims. Insurers will be tempted, in some cases, to reduce commissions on very high-risk (and thus low-margin) cases to contain costs, or might consider those cases as requiring more broker effort (since high-risk populations might need more help with plan management), and thus a larger payment to the broker is warranted. These are difficult considerations to balance. An AI-based system would be capable of learning from historical data how risk factors correlate with necessary broker intervention or completed sales at varying commission rates. For example, it may find that brokers were only able to successfully place coverage for certain high-risk groups when commission was above a certain level, indicating the need for more incentive in those situations, whereas for low-risk, easily insurable groups, standard commissions were sufficient and led to successful conclusions.

The type of healthcare plan being insured also influences commission considerations. Health insurance policies come in dozens of types – health maintenance organizations (HMOs), preferred provider organizations (PPOs), high-deductible health plans combined with savings accounts, comprehensive policies with large networks, limited coverage policies, and dozens more. Certain plans are simple to describe and manage, while others are complex and may require more of the broker's time to inform the client and the workers regarding how the plan works and to iron out issues that arise due to the complexity of the plan. For instance, an HDHP paired with a Health Savings Account might be unfamiliar to some employees, requiring educational seminars and follow-up services on the part of the broker, which would arguably justify a higher commission than a less complicated plan. Plans with narrow provider networks or more aggressive cost control elements might also generate more questions or complaints that must be managed by the broker. Additionally, new or innovative insurance products can mean that the brokers will need to spend time familiarizing themselves with the offering and marketing it. Such plan characteristic differences are rarely explicitly addressed in traditional commission structures—every single plan sold by an insurer can have the same commission rate. However, insurers do occasionally provide special commission promotions for specific products that they want to push (e.g., a carrier might temporarily increase commissions on a new plan type in order to get brokers familiar with and selling it). This demonstrates that plan attributes do have a little bit of applicability to commission in the real world, albeit in an unrefined way. The proposed AI model would consider plan features systematically as a component of the input to determine if the character of a particular plan warrants an exception to the default commission.

The most dynamic variable, perhaps, is the broker's previous performance. In a conventional setting, while base commission levels are usually the same for all brokers for a given product (to avoid preferential treatment and because state insurance regulations often require filing of commission schedules), insurers differentiate brokers through the utilization of the aforementioned bonus or override programs. A broker who retains customers year after year (high persistency or renewal rate), who adds a large amount of new insureds, or who has a book of business with favorable loss ratios (i.e., their clients, on average, cost the insurer less in claims than expected) is more valuable to an insurer's business. These brokers can be indirectly compensated with higher pay or, in some cases, given more leeway when negotiating large accounts. On the other hand, under-performing brokers (e.g., with many clients that switch insurers too frequently or whose groups have unusually high claims) cannot be penalized via reduced base commissions (since they are fixed), but can merely be deprived of additional bonuses and even lose their status with the insurer. In an AI-based commission model, broker performance metrics can directly influence the commission recommendation for each new policy. The model can increase the commission offer as a reward and to retain the broker's interest in placing business with that insurer (basically like providing them with a better deal) if the broker has a good track record. Conversely, if the broker's track record is poor, the model might suggest a standard commission or only a modest increase even for a hard sell, perhaps incorporating the fact that steering valuable clients to top-performing brokers creates more favorable outcomes. Performance metrics can include a wide range of data: policies sold in the previous year, policy retention rate, average client account growth, client satisfaction ratings, compliance with administrative protocols, and more. Based on learning from data, the AI system can deduce which performance metrics are the best predictors of future success and incorporate them into the commission-setting logic.

Apart from these main factors, there are other considerations that sometimes come into commission decisions. Competition in the market is one: when there are many insurers vying for brokers' attention in a line of business, some will increase commissions to encourage brokers to prioritize their products. While the range of our model is internal decision-making within a single insurer, it could be broadened by placing data on rivals' commission levels as an input in order to ensure the proposed commission will enable the insurer to be competitive in obtaining broker business for a particular product or market. Additionally, regulatory constraints can limit or set commissions in some cases (e.g., Medicare Advantage plans have maximum broker commissions set by CMS), but our model primarily addresses the commercial and employer markets where such limits are not uniformly across-the-board and there is more flexibility for the insurer to set commissions. In creating an AI-based solution, these outside variables would be part of the context that data the model is trained on is aware of, even if not explicitly expressed as variables.

A broker's fee made from brokering a health insurance policy is a function of a vector of factors relating to the policy itself, the type of client, the characteristics of the product, along with attributes of the broker himself. Traditional systems capture these factors only to a limited degree—largely through crude segmentation (small group vs. large group, new sale vs. renewal, etc.) and through after-the-fact bonuses. This is an opportunity for more nuance. Our proposed two-stage AI design aims to formalize these influences in a systematic way at the decision point of calculating commission on each policy and thereby operationalize what until now has been a composite of rigid schedules and subjective judgment.

3 Two-Stage AI Framework Architecture

The proposed framework is defined in two successive stages, both of which are motivated by AI techniques, which combined give rise to the ultimate suggestion of commission for a particular insurance policy and broker. Literally, the procedure is similar to the concept of a pipeline: Stage 1 feeds in data on insurance policy and customer and outputs a base commission figure, and Stage 2 processes the base value (and some other broker-specific information) and outputs a revised, final figure for the commission. This design expressly mirrors the practical notion of having a base commission and adjusting performance-wise thereafter but integrates these phases into a data-driven modeling framework.

Stage 1 has the framework focusing on policy-oriented analysis. All the relevant inputs which define the insurance policy and the client environment are fed into an AI model. These inputs are like, as

Healthcare Plan Features Affecting Broker Commissions

Plan Feature	Broker Effort	Typical Handling	AI Proposal
High-Deductible + HSA	Extra client education	Standard rate	Slight commission uplift
Narrow Provider Network	Manage complaints	Occasional bonus	Modest adjustment
New Product Launch	Training and promotion	Temporary bonus	Dynamic scaling
Complex Benefits	Frequent troubleshooting	Ignored	Scored adjustment

Table 3. Plan-specific characteristics affecting broker compensation strategies.

Broker Performance Metrics Impacting Commissions

Performance Metric	Description	Impact	Measurement
Renewal Rate	Client retention percentage	Higher bonus payouts	Annual policy reports
Loss Ratio	Claims vs premiums	Preferential commissions	Underwriting review
New Business Volume	New clients added yearly	Bonus eligibility	Sales data analysis
Client Satisfaction	Feedback scores or complaints	Adjusted incentives	Survey programs

Table 4. Broker performance indicators linked to commission adjustments.

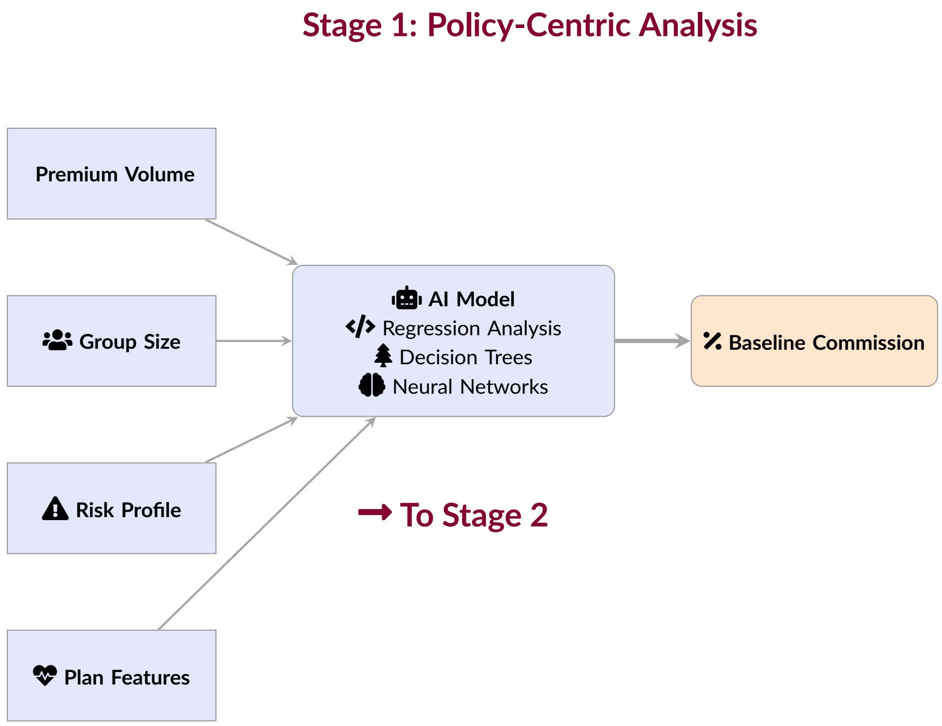

Figure 3. Stage 1 of the proposed two-stage AI framework establishes a baseline commission by analyzing policy and client characteristics independent of broker-specific factors.

discussed in the background, premium volume or policy size, client's risk profile, and key features of the healthcare plan. The Stage 1 model can be compared to a prediction engine that estimates the amount of commission it would be suitable to consider with regard to the effort and the complexity involved based on these inputs alone. In practice, it tries to answer: "For this policy and this client, how much should we compensate a broker on this case, before considering who the broker is?" The Stage 1 output is a benchmark commission recommendation. This could be expressed in absolute dollar terms (e.g., an annual charge or total amount of commission) or relative terms (e.g., a proportion of premium). The internal workings of this model are likely to employ machine learning techniques such as gradient boosting, decision

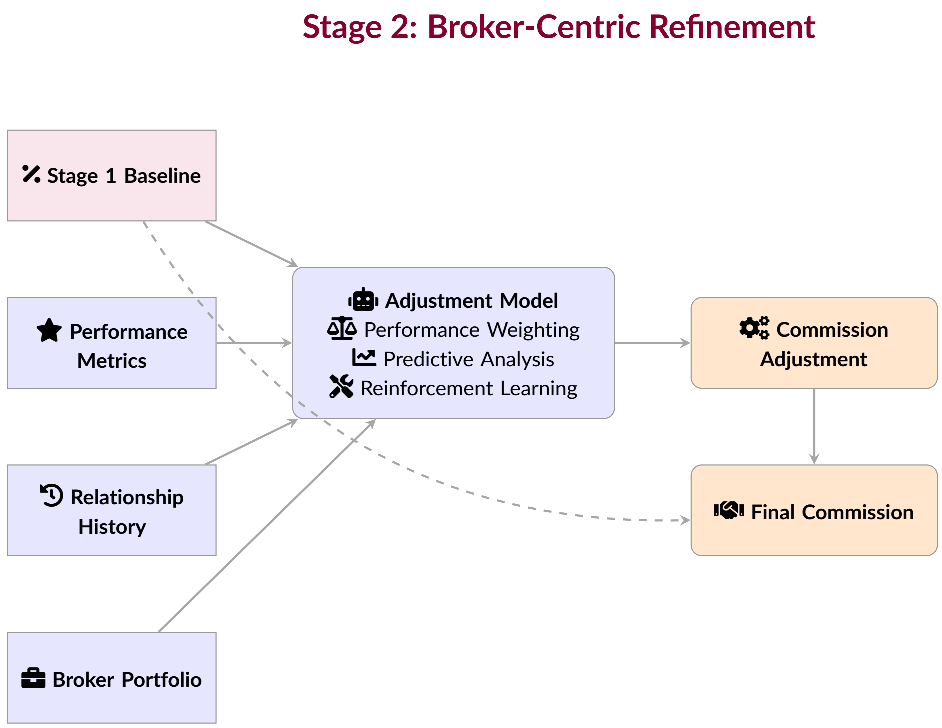

Figure 4. Stage 2 of the proposed two-stage AI framework adjusts the baseline commission from Stage 1 by incorporating broker-specific factors to produce the final personalized commission recommendation.

trees, or neural networks trained on past records of similar policies and outcomes. The training objective might be to predict a commission leading to effective sales and renewals, the level of which is set high enough to elicit broker activity but not high enough to unnecessarily reduce the insurer's margins. As Stage 1 is independent of any broker-specific information, it determines the commission determination as if it were setting an across-the-board rate for any broker depending on the complexity and value of the case.

This stage captures the idea of a data-driven base commission schedule that reacts to case features in much greater detail than one-size-fits-all schedules. Moving to Stage 2, the design incorporates the broker-specific refinement. Stage 2 input includes the Stage 1 baseline commission plus a set of broker-specific performance measurements. The second AI program adjusts the baseline commission higher or lower (or keeps it unchanged) to create an end outcome, separately tailored commission recommendation for said broker in question. Stage 2 serves the purpose of making the payout personalized, i.e., further enhancing Stage 1's one-size-fits-all baseline to one-size-fits-one. The form of AI methodology used at Stage 2 could also differ from Stage 1 because the task is different: Stage 2 could be framed as predictive model (predicting what adjustment leads to favorable broker behavior or outcomes) or even as an optimization or decision model (choosing an adjustment that optimizes some measure of performance). As an example, a supervised learning model in Stage 2 would be trained on historical records where each of various brokers were given varying commissions and then assessing follow-up performance (did they create more business, did the client keep, etc.), basically determining how much influence a given extra commission makes when transacting with brokers of varying quality. Alternatively, one might imagine a reinforcement learning solution in which the model "learns" either through simulation or through repeated data how to calculate commissions to maximize some long-term reward (e.g., broker loyalty or total profitability of the accounts they handle). In either case, Stage 2 generates a commission adjustment which, when used to modify the Stage 1 baseline, gives the final commission rate. In practice, Stage 2 can return a multiplier (for example, 1.10 meaning give a 10% increase for this broker) or an additive figure (for example, add $1,000) or even a fresh rate percentage.

The exact interpretation is left to the implementation; the idea of the framework is indifferent to that point. The two stages are coalesced in such a way that Stage 2 does not override the behavior of Stage 1 but enhances it. Stage 1 ensures the fundamental features of the case are remunerated appropriately, and Stage 2 ensures the exact broker's condition is accounted for. This principle of separation of concerns has various advantages. It is firstly straightforward and transparent: stakeholders are able to identify there is a commission basis of the policy as a standard and then a divergence as a variation. If a broker or manager were to inquire why they were given a specific commission, the answer can be two-fold ("The basis of this kind of policy is X, and you were given an extra Y based on your performance metrics."). It is easier to justify than a single black-box model that produces a number with all considerations at once. Second, modular design can facilitate the system's flexibility and maintenance. If, for instance, the insurer requires adjusting the rewarding of broker performance (Stage 2 logic) without altering the base commission logic, they can retrain or change Stage 2 independently without altering the input and output interfaces.

Similarly, if the conditions in the market shift such that total commissions need to be higher or lower, Stage 1's training or parameters can be adjusted without necessarily adjusting how brokers are differentiated in Stage 2. All stages can subsequently be adjusted on their own cycle, even by separate teams handling separate data sets. Another benefit of the two-stage model is in the model training and data requirements. Stage 1 primarily requires data about policies: their nature, what commission was offered, and what ultimately occurred to those policies (e.g., was the sale successful, did the client stay, did the broker make the effort required). This set of information may encompass a number of brokers' experience without needing them to be distinct, since Stage 1 is broker-agnostic. That means a larger set of training data for Stage 1, making it more stable. Stage 2 requires data corresponding to broker traits, commission variation, and outcomes. By separating the tasks, independent models can be provided with more precise data, which might facilitate learning more efficiently.

Also, technically, two simpler models can work better and are easier to comprehend than one complex model trying to gobble everything up at once—this is like the divide-and-conquer strategy, or using an ensemble of models where each is an expert in part of the task. Operationally, executing the two-stage model could be as follows. Whenever a new policy sale is being contemplated (e.g., a broker is quoting a plan to a new customer), the system of the insurer would first gather all relevant information about that customer and the plan. Stage 1's AI model, which could be deployed as a real-time prediction service, would be invoked to make the baseline commission recommendation. After this, in real time, the system would build the profile of the liable broker – along with his current performance metrics – and enter this as well as the baseline into the Stage 2 model. Stage 2 would then give back the corrected rate or amount. The insurer would then be able to present this final commission offer to the broker as part of the offer. All this could happen in the background within seconds, so it is possible to apply even in live negotiating or automated quoting situations.

In addition, when such a suggestion is proposed and the outcome (whether the transaction was closed, the client was retained, etc.) is observed later on, that data can loop back into an always-learning system to update the models in Stage 1 and Stage 2. This manner, the system is not static; it can evolve with the market and with clients' and brokers' actions, always refining its grasp of the optimal commission arrangements. In conclusion, the architecture of the two-stage AI framework is structured to systematically integrate the complexity of commission determination by splitting it into two manageable components. Stage 1 sets a fair and case-sensitive base commission, and Stage 2 provides broker-specific fairness and strategic incentive alignment. The next section will provide a more in-depth analysis of the mechanisms involved in each phase, the types of models and analyses which could be used, and how the variables like policy size, risk profile, broker performance, and plan features actually function within this system.

4 Framework Components

4.1 Stage 1: Policy-Based Commission Baseline

Stage 1 of the model is concerned with producing a commission recommendation that is solely determined by the features of the insurance policy and the customer's situation and irrespective of the broker handling it. The outcome of this stage is the baseline commission, the starting point before any broker-specific adjustments. In practical terms, this stage answers the question: "What commission would we offer for this policy if we were dealing with an average broker under normal circumstances?"

By isolating policy factors, the model ensures that the underlying needs of the case are met in the commission offer. We can state the Stage 1 model as a function:

$$C_{\text{base}} = f(P, R, L, Z),$$

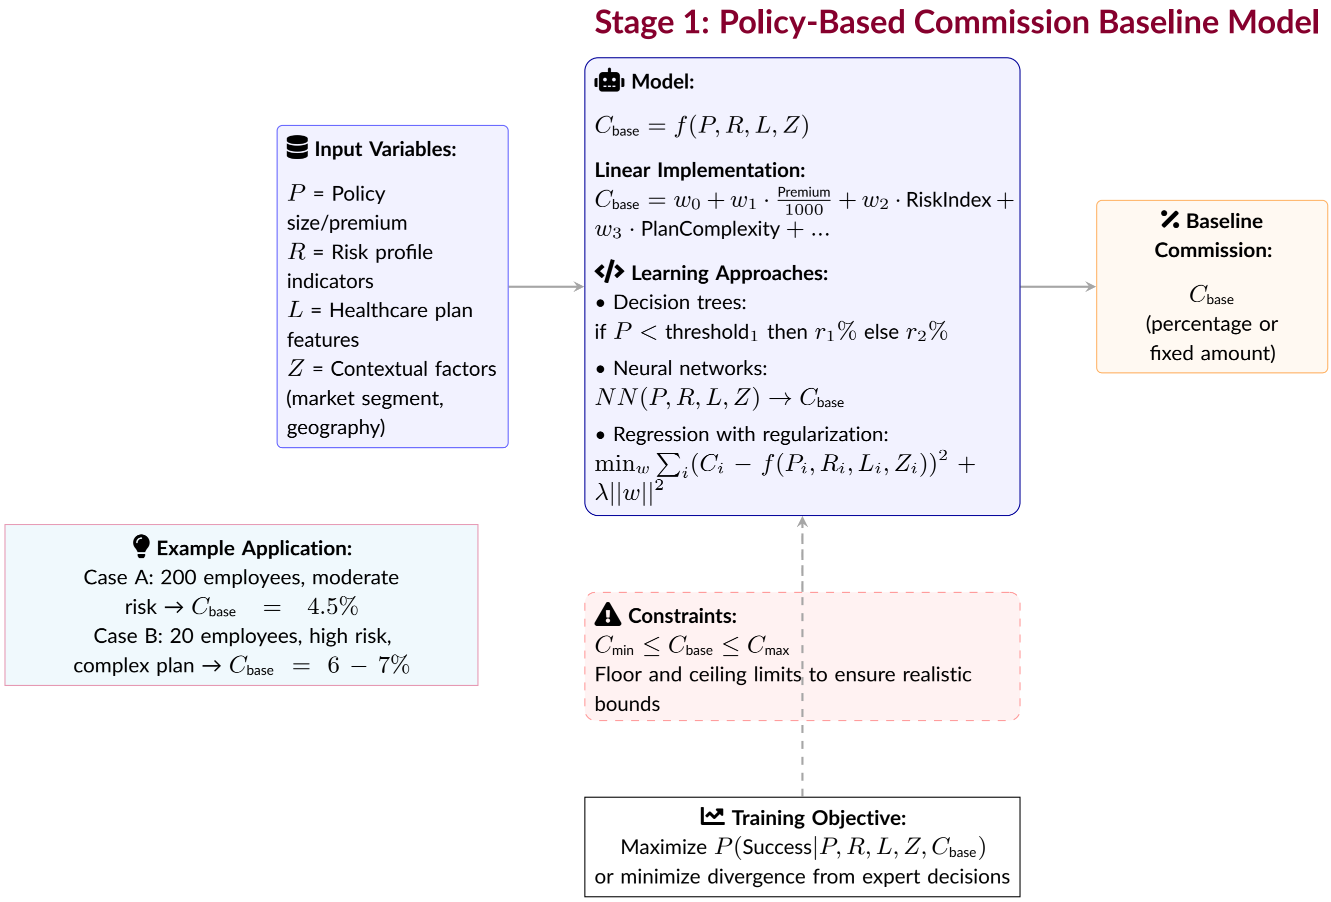

Figure 5. Implementation of Stage 1: Policy-Based Commission Baseline Model. This stage employs functions to derive a baseline commission from policy characteristics, using various machine learning approaches constrained within realistic bounds.

where C_{base} is the baseline commission (perhaps expressed as a percentage of premium or a fixed amount), and the inputs on the right side include P (policy size or premium), R (risk profile indicators of the client or group), L (features of the healthcare plan or product line), and Z (any other relevant contextual factors such as geographic region or market segment). The function f encapsulates the learned relationship between these inputs and the appropriate commission level. In a simple linear form, one might imagine something like:

$$C_{\text{base}} = w_0 + w_1 \cdot \frac{\text{Premium}}{1000} + w_2 \cdot \text{RiskIndex} + w_3 \cdot \text{PlanComplexity} + \ldots$$

where $w_0, w_1, \ldots$ are trained weights from evidence. A linear model is likely to be too simplistic in practice, but it illustrates how different factors can contribute additively to the commission. More sophisticated uses of f could be nonlinear, e.g., with decision trees that apply different percentage commissions based on different premium ranges or neural networks that involve risk and plan feature interactions.

Stage 1 model training would take historical data on previously sold policies. Training examples would include ideally the policy characteristics (size, risk profile, plan type, etc.), the commission provided (or paid) for the policy, and an outcome measure such as whether the sale was closed or how satisfied the broker was (if this data is available). The model can be trained to predict the commission level that will most likely result in good outcomes. In the absence of any particular outcome variable, the model can simply be fit to mimic historical commission decisions, in effect learning the implicit guidelines that have been used historically by the insurer's underwriting and sales staff (which themselves may have been based on experience and intuition of those determinants). Yet the true power of an AI solution lies in the fact that it will identify patterns not captured by linear rules. For instance, the model might conclude that when a client is in a high-risk business (i.e., a business that has been associated with high medical claims) and the group size is medium, commissions slightly above the standard percentage were previously associated with winning the business. The model might then suggest a more liberal baseline commission in similar future cases, even if no hard-and-fast rule previously applied to do so, effectively codifying the intuition that tough cases need a bit more broker incentive.

A crucial aspect of Stage 1 is maintaining the baseline commission within realistic ranges. Model f can be designed with regularization or constraints so that it will not suggest ridiculous values that are evidently unacceptable (e.g., a 50% of premium commission, or an almost $0 commission for a very difficult case). Insurers can realistically apply a floor and ceiling to commissions. The AI model would optimize the commission between these two. If using a machine learning solution, one could enforce these bounds by clipping model outputs or introducing penalty terms during training if the output is outside a reasonable range. The result is that C_{base} will tend to be a plausible percentage like "4.7% of premium" or an amount within the order of what human decision-makers would consider believable, but with more sophistication.

To illustrate how Stage 1 works, let us take a hypothetical example. Suppose that an insurer is evaluating two prospective client cases: Case A is a medium-sized firm with 200 employees in a moderate-risk class, seeking to purchase a standard PPO health plan. Case B is a small firm with 20 employees in a high-risk industry (perhaps one with dangerous work or older populations), looking at a broad health plan with multiple supplemental benefits. Both cases can be quoted, say, a 5% of premium commission by default under a generic commission schedule. But the Stage 1 model can differentiate between them. For Case A, with the larger size, it can determine that a slightly lower percentage (perhaps 4.5%) is sufficient as a minimum, as the premium volume is large and the plan is standard (broker effort per employee is less). For Case B, the model might recommend a higher floor (possibly 6% or 7%) because the group is small (needing more incentive per premium dollar to be profitable for the broker's time) and the client risk profile and plan complexity suggest the broker will have to perform extra work. These amounts are taken from the function f that has learned from similar past instances in which perhaps only small high-risk groups attained coverage when brokers were assured of more-than-normal commissions. In generating such differentiated baselines, Stage 1 explicitly controls for variables that, if left unaddressed, can lead to misaligned incentives (e.g., a broker possibly neglecting a difficult small account if the commission is not commensurate with the effort).

It should be mentioned that the Stage 1 output is not a be-all and end-all; it is the starting point that assumes an average broker. But even in isolation, it could be justified as an excellent improvement over flat commission rules, with a form of risk-adjusted, effort-adjusted compensation. It also sets the stage for Stage 2, by establishing a sensible starting point that can be refined to the individual broker.

4.2 Stage 2: Broker Performance Adjustment

Stage 2 further enhances Stage 1's baseline commission by incorporating broker-specific information into the decision. The goal of this stage is to calibrate the commission so that it is more indicative of the broker's historical performance and behavior. Basically, while Stage 1 answered "how much commission is typical for this case overall," Stage 2 answers "how should that commission be adjusted (if at all) for the particular broker who will handle this case?"

We denote the output from Stage 2 as C_{final}, the final commission recommendation. Conceptually, Stage 2 applies a function g to the base commission and a set of broker performance factors. We can write it as:

$$C_{\text{final}} = g(C_{\text{base}}, B_1, B_2, \ldots, B_m),$$

where $B_1, B_2, \ldots, B_m$ are various quantitative metrics of the broker's performance and characteristics. These would be, e.g., B_1 for the broker retention rate of clients (percentage retained year over year), B_2 for the rate of growth in the book of business of the broker (the amount of new premium they brought in versus their current book for the prior period), B_3 for the average loss ratio of the broker's clients (as a measure of how profitable the business they generate is for the insurer), B_4 for a quality or compliance measure (indicating issues such as whether the broker correctly submits business and follows processes), and B_5 for the tenure or level of experience of the broker. Also, B_5 can represent the experience or seniority of the broker.

Function g can be designed in numerous ways. One simple method is to have g produce a multiplicative factor on the base commission. For instance:

$$C_{\text{final}} = C_{\text{base}} \times (1 + \Delta),$$

where $\Delta = h(B_1, B_2, \ldots, B_m)$ is a percentage adjustment derived from the broker's metrics. If h

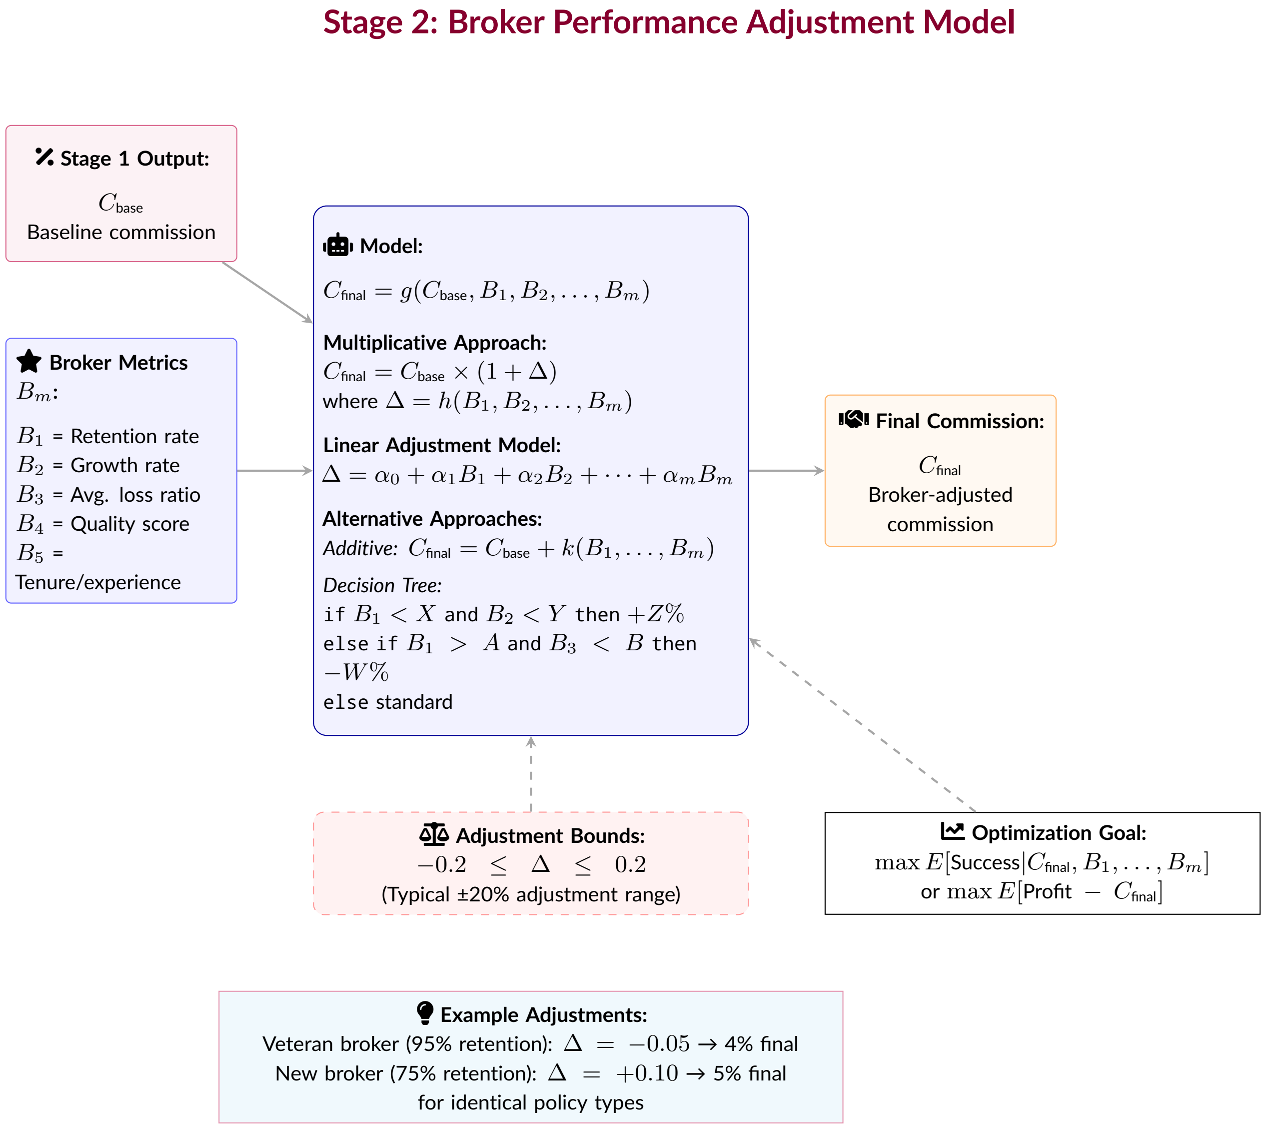

Figure 6. Implementation of Stage 2: Broker Performance Adjustment Model. This stage applies mathematical functions to adjust the baseline commission based on broker-specific metrics, using various adjustment approaches constrained within practical bounds.

outputs 0.10, that means a +10% adjustment, so a baseline commission of \$5,000 would become \$5,500. If h outputs -0.05, that implies a 5% reduction, turning \$5,000 into \$4,750. An alternative formulation is additive:

$$C_{\text{final}} = C_{\text{base}} + k(B_1, \ldots, B_m),$$

where $k(\cdot)$ yields a monetary amount to add or subtract. The multiplicative model tends to be more natural when operating in commission percentages, while the additive model would be more natural if working in flat dollars (appending a flat \$500 fee for an excellent broker to any deal). For the sake of argument, we can consider the multiplicative model; it has the intrinsic property of keeping proportional relationships (the top broker gets, say, 10% more than whatever would have been the baseline for any given scenario).

The AI aspect of Stage 2 is concerned with deciding on what type of adjustment Δ (or the function h in the above example) is suitable. One way to get at this is to look at results that are important to the insurer and then see how they correlate with commissions and broker metrics. As an illustration, consider if the past experience shows that high retention brokers (brokers retaining clients year on year) would do well even without extra commission incentives; they might not need a high increase in the commission to perform well since their incentive could be service commitment or existing revenue source. In contrast, newer brokers or those with lower retention can dramatically enhance their effort or focus for a client if the commission is greater. Such a model could go through the information and learn to recognize an optimum pattern, perhaps something like this: "For a broker retaining less than 80%, the extra 15% of baseline commission was typically needed to earn the same results in sales and retention as the higher-retaining brokers.". If retention rate is more than 95%, commission might even be 5% lower than baseline and the broker would still do well. Such a pattern, once learned, would be encoded in the function g.

As an actual definition, consider a simplified model where g is linear in broker features for a multiplicative modification:

$$\Delta = \alpha_0 + \alpha_1 B_1 + \alpha_2 B_2 + \cdots + \alpha_m B_m.$$

Here, α_0 might be zero (i.e., no correction for an average broker), and α_1, α_2, etc., are learned coefficients by the algorithm to fit. If B_1 is retention rate (as a decimal fraction), a negative α_1 would mean greater retention results in a negative Δ (a decrease in commission, i.e., the broker does not need so high a commission). Conversely, a positive α_2 under the growth rate criterion would mean that brokers actually growing their book would be rewarded with a higher commission (e.g., to incentivize and encourage their book building). Linear model is easy to understand but might not be capable of handling nuances – e.g., a broker might need to meet many requirements (high growth and high quality scores) before a rise is warranted, which is a nonlinear interaction.

So increasingly powerful models for g can be used, such as decision trees to produce broker segments (e.g., "if retention $<$ X and growth $<$ Y, then +Z% commission; else if retention $>$ A and loss ratio $<$ B, then -W% commission; otherwise default"). A rule-based meaning might be something like: top-performing brokers (high retention, profitable transaction) might dip slightly from baseline or no adjustment, mid-ranking brokers get baseline, and the lowest performers get increased commission to encourage them or to pay for extra advice they might require in order to consummate the sale. The AI system does not necessarily have to be manually coded with these rules – it would learn them from results (such as which deals were closed, which clients retained, etc., under what broker and commission conditions).

It should be noted that a "reduced commission" for an elite broker is not necessarily paying them less absolutely than a less capable broker. Because Stage 1 is already causing commissions to rise proportionally with the case, an top-shelf broker is often dealing with larger or more complex cases anyway and would make more in absolute dollars on them even if this weren't being done. Stage 2 just nudges at the margins, perhaps only not paying too much when unnecessary and diverting it somewhere else where it will do more good. In practice, an insurer could write that Stage 2 may differ only within a given range (e.g., $\pm$20% of base case) to achieve equity and preclude gross inequities.

A circumstance under which Stage 2 would be particularly valuable is if an insurer has some old-timers and some new blood brokers. The veteran might have many established customers and might be working on a new customer as an accommodation or favor; they might do it even if the commission is slightly lower because they don't want to lose the relationship with the insurer or the customer. A new broker trying to establish their reputation might seek out higher-commission deals and might demand that extra cash in order to spend the time that a complex sale demands. By recognizing such differences, Stage 2 would be in a position to offer, for the same policy type, say a 4% commission to the veteran (with confidence they will do all right anyway) but a 5% commission to the rookie (to get their attention and work). Such segmented treatment can generate overall better outcomes for the insurer's sales growth and customer satisfaction, since each broker is being suitably encouraged.

From a data perspective, training Stage 2's model would involve examining historical changes in broker performance and how commissions had correlated with success for different broker profiles. If the insurer has experimented with different commission levels or has seen natural experiments (e.g., changes in commission schedules) and observed broker behavior, those would be gold to train on. Even without direct experiments, nature can be a source of variation (different lines of business may have had different commission rates) that can be informative for learning signals. The model would attempt to optimize something similar to the insurer's policy success expectation (such as sale obtained and client retained for at least a period of X years) or insurer profit (with commission as an expense and earned premiums as income) by adjusting commissions. This effectively turns Stage 2 into a mini-optimization engine upon the base provided by Stage 1.

At deployment, following a Stage 2 recommendation of a new commission, the insurer can communicate this as proposed commission to the broker. It can be phrased in the context of an incentive program or simply as the commission on that case. Brokers might ultimately come to understand that their performance influences their compensation offer. This can have the added advantage: an incentive to brokers to improve those performance levels (e.g., retention and expansion) because they realize it will directly affect their revenue potential. However, the system's ultimate purpose is not to establish a reward or punitive system in an arbitrary fashion but to enable information to direct an optimal allocation of commission dollars. That is, it ensures that the insurer is investing the commission budget where it brings the highest return in terms of business objectives.

Stage 2, combined with Stage 1, gives us the full framework. By the end of it, we have a terminal value of commission C_{final} which has catered to case-specific needs as well as broker-specific circumstance. The final counsel is now ready to be added to the company procedure of determining the price on the insurance offer and paying for the broker sale.

5 Analytical Formulation and Application

To formalize the two-stage model, in short, we can consider the commission determination as a function that ultimately depends on both policy-level characteristics as well as broker-level characteristics. Let X be the policy and client characteristics (premium, risk factors, plan details, etc.) and B be the set of broker characteristics (performance measures and corresponding metrics). The two-stage method, in essence, defines the commission function compositionally:

$$C_{\text{final}}(X, B) = g(f(X), B).$$

Here, $f(X)$ corresponds to the Stage 1 model's output C_{base}, and $g(Y, B)$ (with $Y = f(X)$) corresponds to the Stage 2 adjustment logic producing C_{final}. If we use a multiplicative adjustment form for simplicity, this can be expanded as:

$$C_{\text{final}}(X, B) = f(X) \times [1 + h(B)],$$

where $h(B)$ yields the adjustment factor based on broker features (so $h(B) = 0.10$ would mean a +10% adjustment, etc.). This equation encapsulates the heart of the framework: the commission is first determined by the characteristics of the policy X, and then scaled up or down according to the characteristics of the broker B.

It is instructive to note that if we had chosen not to use a two-stage approach, we would attempt to learn a single function $F(X, B)$ mapping all inputs directly to C_{final}. While that is possible, the two-stage structure $g(f(X), B)$ provides interpretability and mirrors domain practice (base plus adjustments). Moreover, it allows the insurer to interpret $f(X)$ as the "standard" commission for a case and g or $h(B)$ as the deviation for a particular broker. In a fully learned single function, these distinctions blur, which could make it harder to justify or adjust the commission policy.

The model presented here can similarly be used from an optimization perspective. The insurer basically wants to choose C_{final} which maximizes some outcome, e.g., the probability of a successful sale and policy maintenance, minus the cost of commission. The AI functions f and g are set up to approximate the optimal policy for computing commissions in various situations. If we allow $U(\cdot)$ to be an insurer's objective function (e.g., something that increases with the probability of writing profitable business and decreases with higher commission cost), then for each scenario (X, B) the chosen C_{final} is intended to mimic:

$$\arg\max_{C} U(\text{outcome} \mid X, B, C) - \lambda[C],$$

where λ is a cost weighting for the commission cost. The models f and g learn this indirectly by observing what decisions led to good results in the past. Although we do not carry out this optimization here explicitly in our conceptual paper, establishing it in this way helps us to see that the AI-based framework is attempting to allocate commission resources optimally for the outcomes from the data.

To give the operation of two-stage method some more real-world example, let us consider an example scenario. Let there be an insurance company dealing with a prospective customer, and the customer is a small technology startup company (let's refer to the customer as Client Z) with 50 employees. The client has quite a healthy and young staff, but they are in a stressful working environment that can result in average health claims. The client is considering a fairly large PPO medical plan with mental health and broad provider network (which is a costly plan). From these observations, the carrier constructs the features:

- P: Policy size = 50 employees (premium, e.g., \$400 per employee per month, which is \$20,000 per month or \$240,000 yearly premium). - R: Risk profile = moderate (perhaps an actuarial risk index of 1.1,

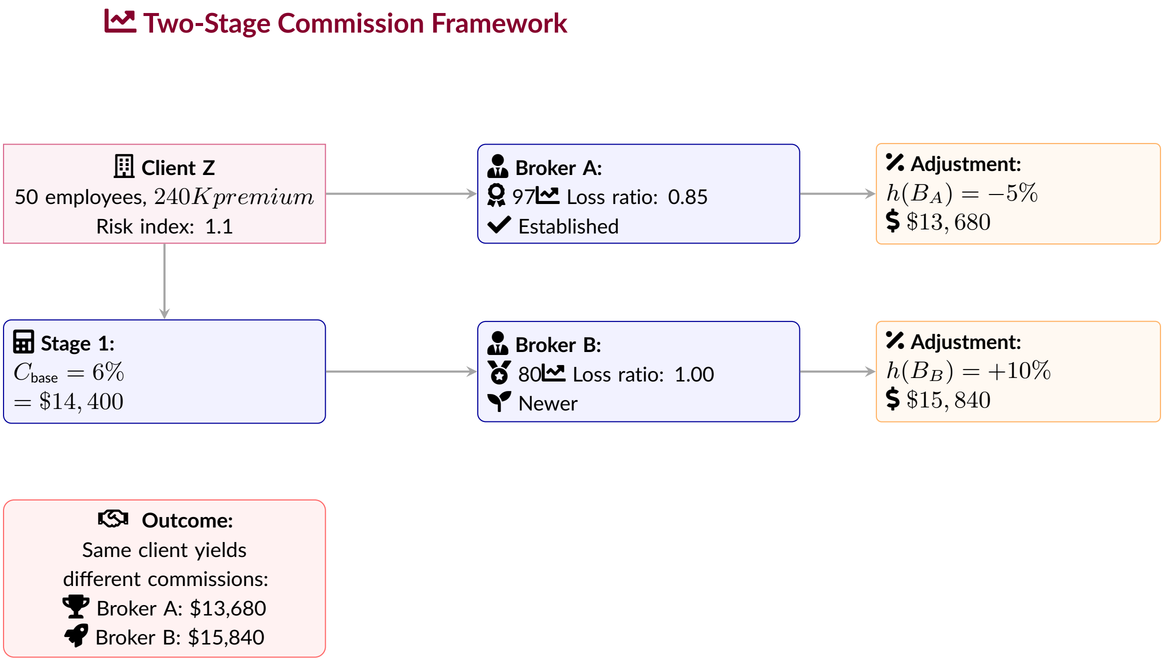

Figure 7. Two-stage commission framework highlighting how broker characteristics affect final commission rates.

slightly above average, to reflect some stress-related claims potential but a young population). - L: Plan characteristics = high coverage breadth and high network flexibility (we can quantify this qualitatively as a "complexity" score of perhaps 8 out of 10, since it's a full-coverage plan which requires description). Plugging this $X = (P, R, L)$ into the Stage 1 model $f(X)$, supposing that the AI has learned that in instances with groups of around 50 lives and fairly higher-than-average risk and a sophisticated plan, the commission ought to be at the high end of the norm to induce broker involvement, a sample output would be:

$$f(X) = 0.06 \text{ (i.e., 6\% of premium)}.$$

This would be equivalent to $C_{\text{base}} = 0.06 \times \$240,000 = \$14,400$ as the base annual commission for this case. For comparison purposes, if the insurer had a general rule of 5% for groups of this size, our model is suggesting a bit more, likely for the complexity and moderate risk warranting more broker effort.

And now, comparing two different brokers who might be awarded this client:

Broker A is an experienced broker with a stellar record: 97% retention of clients, consistently growing her business, and with an average loss ratio on her accounts of 0.85 (i.e., claims are 85% of premiums, which is relatively profitable business on average). Broker A often does not need to go after new business aggressively; she often gets referrals due to her reputation.

Broker B is a new broker with moderate success but not dazzling: 80% retention (a handful of clients have left in the past couple of years), slow growth mostly due to sheer grind prospecting, and an average loss ratio of 1.00 (his clients' premiums roughly equal claims, on average). Broker B is looking to make a name for himself and tends to chase opportunities that will reward him nicely, since he's just building his income stream.

We are inputting Broker A's figures B_A and Broker B's figures B_B into Stage 2 model. Assume that the model $h(B)$ (under the multiplicative adjustment view) returns us with:
- For Broker A: $h(B_A) = -0.05$ (i.e., -5% adjustment).
- For Broker B: $h(B_B) = +0.10$ (i.e., +10% adjustment).

These figures fit into an explanation of how Broker A would welcome having a bit lower in commission since she will either close and service the transaction efficiently anyway (perhaps because she already has the client's acquaintance or because she possesses better ability), while Broker B would need to be pushed in order to give the right amount of attention and to close the deal efficiently.

Now adjusting these: - For Broker A: $C_{\text{final}} = \$14,400 \times (1 - 0.05) = \$14,400 \times 0.95 = \$13,680$. This is an effective commission rate of around 5.7% of premium for Broker A. - For Broker B: $C_{\text{final}} = \$14,400 \times (1 +$

$0.10) = \$14,400 \times 1.10 = \$15,840$. This is an effective commission rate of around 6.6% of premium for Broker B.

Therefore, in the same exact client scenario, Broker B would be rewarded with $15,840 commission while Broker A would receive $13,680. Broker B's offer is more in absolute and relative terms to incentivize a broker who otherwise might not be as enthusiastic or productive on this tough sale. Broker A's bid is slightly under the benchmark, no doubt because she might not require so high a commission to be able to return a good profit. To repeat, Broker A is by no means "underpaid" in this case – even at 5.7%, she's taking home a staggering $13.7k on one account alone, and her healthy portfolio quite likely earns her cash from many other accounts. Broker B is being given another $1,440 on top of the base to make this deal his top priority, if it means anything to him.

What can we anticipate as a result? If the learned patterns of the model are correct, Broker B's extra effort (because of the increased commission) may make it more likely that Client Z signs up and stays a happy customer, while Broker A would probably have sold Client Z successfully even at the reduced commission. In most cases, the insurer is really investing its commission expense where it will do the most good – paying a bit more only when needed in order to acquire or maintain business, and paying a bit less where the business likely would stay even without added incentives. This can mean a cheaper growth plan.

If we suppose a very large client (5,000 employees, highly desirable risk profile) is on the table. Stage 1 might give a baseline commission of perhaps 3% (due to the huge premium amount, even a small percentage commands a massive dollar commission). And if Star Broker A can do it, Stage 2 might drop it to 2.8%, netting, let's say, $420,000 instead of $450,000—Star Broker A might not even notice the difference relative to the size of the deal, and she already has a lot riding on getting this kind of business. If Broker B obtains it, Stage 2 will take it to the bottom line 3% or even slightly higher to 3.3%, paying him maybe $495,000, which could be a huge chance for him. In either case, the margin is small in relation to the magnitude of the deal, but amplified by performance.

In an applied application, the actual dollar amounts and percentage points would be found from large pools of data and strictly audited but the qualitative outcome should be: commissions that adjust according to the circumstance of the sale and behavior of the broker, not an ad hoc instrument.

6 Discussion

The use of an AI-based two-stage commission scheme has several consequences for the various parties involved in the health insurance industry. A direct consequence is that there could be more efficient incentives distribution. By calibrating commissions to the detail of individual sales and of each broker, the same insurer can achieve comparable or even superior business outcomes (in policy sales, client retention, and profitability) without paying too much in commission expenses. That is, the insurer is insuring most where they are needed most. Eventually, this cost-effectiveness would manifest in the form of less total administrative cost or the ability to invest those savings in such items as customer care or product development. Brokers, on their part, would be fulfilling a compensation formula that is responsive to their own performance. High-end brokers see this as a nod to their efficiency—if they are able to regularly get the business closed and keep clients happy even at lower rates of commission, it is a sign of their strength in the market, and they still make money by being able to get more volume done. Less skilled or lower-performing brokers would presumably welcome the extra commission incentives that fall their way for difficult cases in the first place, and hopefully utilize them as an incentive to improve their own practices.

The other main point of discussion is how this type of framework might influence broker behavior and market dynamics. If brokers find that particular behaviors or performance measures (e.g., retention rates or customer satisfaction) affect their provision of commission, they will in turn direct more attention to them. This can produce a virtuous cycle: brokers competing to keep customers and deliver good service in order to negotiate better remuneration terms, which ultimately benefits the insurers and customers through more sustainable long-term relationships and better customer experiences. In fact, the system not only reacts to broker performance but can even encourage better performance. But care must be taken to ensure that the metrics used actually measure good outcomes and do not incentivize any counterproductive behavior (e.g., a broker might overemphasize retention at the expense of new business, if only retention is incentivized; a balanced set of metrics in B can prevent such traps).

The design also offers a form of personalization in B2B relationships that is increasingly common in many industries due to AI: just as shoppers see tailored pricing or promotions online, brokers (business partners of insurers) may see tailored commission proposals. This would strengthen alliances, as the broker feels that the insurer is tailoring the offer to them, as opposed to presenting them with a take-it-or-leave-it base rate. It identifies the individual contribution of the broker. From the opposite standpoint, insurers would need to approach this with some care in order to maintain an ingredient of fairness and transparency. Even though we have avoided ethical and regulatory grounds, it is reasonable to say from a business perspective that gross inequality in proposals for commission would need to be justified by clear inequality in situation to avoid resentment by brokers. In the real world, insurers would typically apply this scheme moderately—minutely up and down—such that the overall compensation climate remains level and equable but continues to leverage on having the degree of fine tuning. AI technology is capable of being conservative on adjustments especially for roll-out onset and learning throughout from the optimum sizes of the adjustment that produce payoff without inconvenience.

Technically, an issue and area of contention is data and model management. To employ the two-stage framework, the insurer would need to gather high-quality historical data on policies, broker performance, commissions, and outcomes. Data silos would need to be broken down: often, broker performance data will be in a separate system than policy underwriting data or sales data. Sustaining these collectively for modeling purposes is a not-trivial undertaking. Furthermore, the models themselves need to be monitored over the course of time. If there is a change in the health insurance market (e.g., new legislation in broker compensation or modifications in the broker function as a result of direct-to-consumer online selling), the AI models need to be re-trained or calibrated for them to perform. Fortunately, its modular design is such that one module can be rewritten in isolation—stage 2 can be trained again if behavior shifts at brokers without the need to retrain stage 1 as long as the underlying relationship between policy factors and necessary commission (stage 1) does not change, and so forth [44], [45].

Of further interest are how scalable and generalizable this framework is. While we used it in the context of U.S. health care insurance, the two-stage commission-setting model could be applied to other segments of insurance or even to other markets where there are agents or intermediaries who get paid. For example, property insurance brokers or life insurance agents could also benefit from a more advanced commission structure. The Stage 1 and Stage 2 variables would differ (for life insurance, for example, the client's mortality risk, etc.), but the concept of breaking down the decision into a case-base and an agent-adjustment would be applicable. Such a system could conceivably accommodate future trends in healthcare as well, such as expansion of digital brokerages or broker services via artificial intelligence; even those could be "graded" for performance and given appropriate commission incentives.

In rolling out such a system, insurers could take a staged approach: use the AI system as a guide within, and human judgment at first, compare outcomes, and gradually learn to trust the model. Over time, with growing confidence in the AI recommendations, it could be tasked with determining commissions increasingly independently. This phased-out method would be a gradual way to obtain the benefits of the framework without restricting disruption. Throughout, the overriding guideline needs to be that the AI-driven framework is a tool for enhancing decision-making, and not of replacing strategic management [46], [47]. If used judiciously, it can provide insights (such as which kinds of cases have been under- or over-incentivized in the past) and refine the broker compensation strategy in a manner that is evidence-based and aligned with the insurer's goals.

7 Conclusion

The present paper presented a conceptual framework to determine insurance broker commissions in the U.S. healthcare industry through a two-stage artificial intelligence-based method. By dividing the process of commission determination into a policy stage and a broker stage, the framework systematically combines a full range of variables that influence commission requirements. The first step establishes a base commission that suits the specificities of the health insurance policy and client, considering the policy size, client risk profile, and type of healthcare plan. The second step modifies this offer according to the previous performance of the broker in order to ensure that the final commission will be in line with the capabilities demonstrated by the broker and incentive requirements

Targeting the U.S. context of the healthcare system, we confronted the complexity and diversity in health

Impact of AI-Driven Commission Model on Stakeholders

Party	Benefit	Challenge	Note
Insurers	Cost efficiency	Fairness management	Lower admin costs
High Performers	Recognition, growth	Maintain service quality	Higher volumes achievable
Low Performers	Incentive to improve	Risk of dependency	Targeted support needed
Clients	Better service	Risk of metric imbalance	More sustainable policies

Table 5. Impacts of AI-based commission schemes on different stakeholders in the healthcare insurance system.

Predicted Broker Behavior Under AI Commission Model

Behavior	Motivation	Risk	Management
Focus on Retention	Bonus on renewals	Neglect of new sales	Balanced metrics
Service Quality Push	Satisfaction scores matter	Gaming feedback	Robust surveys
Volume Expansion	Reward for new business	Lower deal quality	Quality checks
Risk Selection	Favor stable groups	Skewed target market	Diversity incentives

Table 6. Expected broker behavioral shifts in response to AI-driven commission adjustments.

insurance distribution. In the field, brokers have been compensated on relatively inflexible structures, but the insurance-selling is full of nuances – smaller or riskier accounts might require more broker effort, and brokers are heterogeneous in efficiency and motivation. The proposed AI system includes these nuances: Stage 1 covers the understanding that different insurance scenarios need different rates of commission, and Stage 2 includes a customized element that rewards or incentivizes brokers in a suitable way. This architecture captures usual industry norms (e.g., base commissions plus bonus incentives) but applies data and machine learning to calibrate those components more accurately than is typical.

One of the advantages of this two-stage approach is its transparency and flexibility. Both stages can be defined and interpreted separately, allowing stakeholders to see the rationale behind a commission rate. The use of AI means that as market conditions change or new information is introduced, the models can learn and revise the commission recommendations, potentially leading to optimally optimized outcomes. If new types of health plans are created or broker behavior changes, the framework is flexible enough to accommodate those changes by amending the Stage 1 or Stage 2 models, respectively. Furthermore, by allocating commission funds in a targeted way – giving more where it adds greater value and less where it is not required – the system could help insurers maintain better control of their distribution costs while still enjoying broker engagement and satisfaction.

This is theory work and presents a model rather than a report of empirical results. The next steps towards application in the real world would involve gathering the data required, training the machine learning algorithms, and verifying the recommendations of the framework against real-world results. Problems such as insurer process integration and transparency to brokers would be important to take-up, but those were beyond the scope of this paper. Still, the model provides a foundation for viewing broker commissions as a variable, data-based solution rather than an inflexible cost.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgement

This work was supported without any funding.

References

[1] S. Mutenga and S. K. Staikouras, “The theory of catastrophe risk financing: A look at the instruments that might transform the insurance industry,” *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 32, no. 2, pp. 222–245, Apr. 16, 2007. DOI: 10.1057/palgrave.gpp.2510127.

[2] D. Oughton, "Expanding tort liability in english law and compulsory insurance for professional risks," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 14, no. 4, pp. 331–346, Oct. 1, 1989. DOI: 10.1057/gpp.1989.30.

[3] P.-R. Chang, J.-L. Peng, and C. K. Fan, "A comparison of bancassurance and traditional insurer sales channels," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 36, no. 1, pp. 76–93, Jan. 6, 2011. DOI: 10.1057/gpp.2010.34.

[4] M. Ogrizek, "The effect of corporate social responsibility on the branding of financial services," *Journal of Financial Services Marketing*, vol. 6, no. 3, pp. 215–228, Mar. 1, 2002. DOI: 10.1057/palgrave.fsm.4770053.

[5] H. Cousy and H. Claassens, "Ex post control of insurance in belgium," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 19, no. 1, pp. 46–59, Jan. 1, 1994. DOI: 10.1057/gpp.1994.5.

[6] F. G. Nogueira, A. F. Lucena, and R. Nogueira, "Sustainable insurance assessment: Towards an integrative model," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 43, no. 2, pp. 275–299, Sep. 14, 2017. DOI: 10.1057/s41288-017-0062-3.

[7] M. Menear, M.-A. Blanchette, O. Demers-Payette, and D. Roy, "A framework for value-creating learning health systems," *Health research policy and systems*, vol. 17, no. 1, pp. 79–79, Aug. 9, 2019. DOI: 10.1186/s12961-019-0477-3.

[8] J. Mellow and R. B. Greifinger, "Successful reentry: The perspective of private correctional health care providers," *Journal of urban health : bulletin of the New York Academy of Medicine*, vol. 84, no. 1, pp. 85–98, Nov. 28, 2006. DOI: 10.1007/s11524-006-9131-9.

[9] H. A. Wondirad, "The impacts of mobile insurance and microfinance institutions (mfis) in kenya," *Journal of Banking and Financial Technology*, vol. 4, no. 1, pp. 95–110, May 30, 2020. DOI: 10.1007/s42786-020-00021-2.

[10] R. Cooper, J. Seddon, and B. V. Vliet, "High-frequency trading and conflict in the financial markets," *Journal of Information Technology*, vol. 32, no. 3, pp. 270–282, Sep. 1, 2017. DOI: 10.1057/s41265-016-0031-5.

[11] J. A. Leppo, "Seventh annual mario s. verani, md memorial lecture: "asnc and integrating cardiac imaging"," *Journal of nuclear cardiology : official publication of the American Society of Nuclear Cardiology*, vol. 16, no. 2, pp. 184–192, Feb. 14, 2009. DOI: 10.1007/s12350-009-9054-4.

[12] G. Benoist, R. Harvey, and F. Mira, "Additional remarks and comments on global pension schemes the french case," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 28, no. 1, pp. 26–38, Jan. 1, 2003. DOI: 10.1111/1468-0440.00203.

[13] C. D. O'Brien, "Measuring the output of life assurance companies," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 16, no. 2, pp. 207–235, Apr. 1, 1991. DOI: 10.1057/gpp.1991.17.

[14] M. S. Wijenaike, "Challenges and prospects of developing in asian countries in the non-life reinsurance markets," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 20, no. 2, pp. 148–165, Apr. 1, 1995. DOI: 10.1057/gpp.1995.12.

[15] A. Bagnasco, N. Dasso, S. Rossi, *et al.*, "Esicm lives 2019 : Berlin, germany. 28 september - 2 october 2019 - 2 abstract: 1) nurses' wellbeing in neonatal and pediatric intensive care units and its relations with perceived nursing sensitive outcomes; 2) how does healthcare quality influence care left undone in neonatal and paediatric intensive care units?" *Intensive care medicine experimental*, vol. 7, no. 55, pp. 55–55, Sep. 27, 2019. DOI: 10.1186/s40635-019-0265-y.

[16] J. D. Arpel, "Environment and insurability the application of risk management principles to running an insurance company," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 20, no. 3, pp. 374–392, Jul. 1, 1995. DOI: 10.1057/gpp.1995.31.

[17] C. J. Guerrini, J. K. Wagner, S. C. Nelson, G. H. Javitt, and A. L. McGuire, "Who's on third? regulation of third-party genetic interpretation services," *Genetics in medicine : official journal of the American College of Medical Genetics*, vol. 22, no. 1, pp. 4–11, Aug. 12, 2019. DOI: 10.1038/s41436-019-0627-6.

[18] K. C. Ferdinand and A. Armani, "Combating racial and ethnic disparities in cardiovascular disease: What is being done?" *Current Cardiovascular Risk Reports*, vol. 1, no. 2, pp. 124–132, May 7, 2007. DOI: 10.1007/s12170-007-0020-x.

[19] V. Bentotahewa, C. Hewage, and J. Williams, "The normative power of the gdpr: A case study of data protection laws of south asian countries," *SN Computer Science*, vol. 3, no. 3, Mar. 7, 2022. DOI: 10.1007/s42979-022-01079-z.

[20] P. Kaur and A. Parashar, "A systematic literature review of blockchain technology for smart villages.," *Archives of computational methods in engineering : state of the art reviews*, vol. 29, no. 4, pp. 1–52, Oct. 27, 2021. DOI: 10.1007/s11831-021-09659-7.

[21] V. D. Lorenzo, "Business ethics: Law as a determinant of business conduct.," *Journal of Business Ethics*, vol. 71, no. 3, pp. 275–299, Nov. 3, 2006. DOI: 10.1007/s10551-006-9139-9.

[22] O. Schöffski and A. G. Wegener, "Risk management and insurance solutions for space and satellite projects," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 24, no. 2, pp. 203–215, Apr. 1, 1999. DOI: 10.1111/1468-0440.00016.

[23] S.-H. Li, C.-Y. Wang, W.-H. Lu, Y.-Y. Lin, and D. C. Yen, "Design and implementation of a telecare information platform," *Journal of medical systems*, vol. 36, no. 3, pp. 1629–1650, Dec. 1, 2010. DOI: 10.1007/s10916-010-9625-6.

[24] A. Shreck, C. Getz, and G. Feenstra, "Social sustainability, farm labor, and organic agriculture:

Findings from an exploratory analysis," *Agriculture and Human Values*, vol. 23, no. 4, pp. 439–449, Nov. 21, 2006. DOI: 10.1007/s10460-006-9016-2.

[25] W. R. Majure, "The antitrust division: Economic issues presented in 2010–2011," *Review of Industrial Organization*, vol. 39, no. 4, pp. 327–334, Nov. 18, 2011. DOI: 10.1007/s11151-011-9324-6.

[26] L. B. DeLind and P. H. Howard, "Safe at any scale? food scares, food regulation, and scaled alternatives," *Agriculture and Human Values*, vol. 25, no. 3, pp. 301–317, Jan. 16, 2008. DOI: 10.1007/s10460-007-9112-y.

[27] D. A. Bailey, "D&o liability in the post-enron era," *International Journal of Disclosure and Governance*, vol. 2, no. 2, pp. 159–176, Jun. 1, 2005. DOI: 10.1057/palgrave.jdg.2040050.

[28] D. Busch and M. B. J. van Rijn, "Towards single supervision and resolution of systemically important non-bank financial institutions in the european union," *European Business Organization Law Review*, vol. 19, no. 2, pp. 301–363, Mar. 28, 2018. DOI: 10.1007/s40804-018-0107-5.

[29] P. Karaca-Mandic, J. M. Abraham, and C. E. Phelps, "How do health insurance loading fees vary by group size?: Implications for healthcare reform.," *International journal of health care finance and economics*, vol. 11, no. 3, pp. 181–207, Aug. 20, 2011. DOI: 10.1007/s10754-011-9096-4.

[30] "Trade in services," *United Kingdom Balance of Payments - The Pink Book*, vol. 2009, no. 1, pp. 42–55, Aug. 14, 2009. DOI: 10.1057/ukbp.2009.4.

[31] W. G. Poortvliet and T. P. Laine, "A global trend: Privatization and reform of social security pension plans.," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 19, no. 3, pp. 257–286, Jul. 1, 1994. DOI: 10.1057/gpp.1994.19.

[32] M. D. M. Davis, A. Schermuly, A. K. J. Smith, and C. Newman, "Diversity via datafication? digital patient records and citizenship for sexuality and gender diverse people," *BioSocieties*, vol. 18, no. 2, pp. 451–472, Apr. 5, 2022. DOI: 10.1057/s41292-022-00277-5.

[33] J.-A. Fiske and A. J. Browne, "Aboriginal citizen, discredited medical subject: Paradoxical constructions of aboriginal women's subjectivity in canadian health care policies," *Policy Sciences*, vol. 39, no. 1, pp. 91–111, May 13, 2006. DOI: 10.1007/s11077-006-9013-8.

[34] M. Rival and R. Major, "What lobbying ethics and what for? the case of french lobbying consulting firms," *Journal of Business Ethics*, vol. 150, no. 1, pp. 99–116, Apr. 5, 2016. DOI: 10.1007/s10551-016-3141-7.

[35] J. E. Thomas, "The role and powers of the chinese insurance regulatory commission in the administration of insurance law in china," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 27, no. 3, pp. 413–434, Jul. 1, 2002. DOI: 10.1111/1468-0440.00182.

[36] C. R. Wentink, S. Raemaekers, and S. R. Bush, "Co-governance and upgrading in the south african small-scale fisheries value chain," *Maritime Studies*, vol. 16, no. 1, pp. 5–, Apr. 5, 2017. DOI: 10.1186/s40152-017-0058-4.

[37] M. Faure, "Alternative compensation mechanisms as remedies for uninsurability of liability," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 29, no. 3, pp. 455–489, Jul. 1, 2004. DOI: 10.1111/j.1468-0440.2004.00297.x.

[38] C. Kretzler and F. Wagner, "The catastrophe risk exchange (catex) in new york: An exchange and electronic processing system for reinsurance business," *The Geneva Papers on Risk and Insurance - Issues and Practice*, vol. 25, no. 1, pp. 136–153, Jan. 1, 2000. DOI: 10.1111/1468-0440.00054.

[39] S. Chai and M. A. Scully, "It's about distributing rather than sharing: Using labor process theory to probe the "sharing" economy," *Journal of Business Ethics*, vol. 159, no. 4, pp. 943–960, Jun. 12, 2019. DOI: 10.1007/s10551-019-04210-y.

[40] M. Alles, S. M. Datar, and J. Friedland, "Governance-linked d&o: Market-based governance: Leveraging d&o insurance to drive corporate governance," *International Journal of Disclosure and Governance*, vol. 3, no. 2, pp. 84–98, Jun. 1, 2006. DOI: 10.1057/palgrave.jdg.2040068.

[41] L. S. Parker, T. Karliychuk, D. Gillies, B. Mintzes, M. Raven, and Q. Grundy, "A health app developer's guide to law and policy: A multi-sector policy analysis.," *BMC medical informatics and decision making*, vol. 17, no. 1, pp. 1–13, Oct. 2, 2017. DOI: 10.1186/s12911-017-0535-0.

[42] K. Morrell and I. Clark, "Private equity and the public good," *Journal of Business Ethics*, vol. 96, no. 2, pp. 249–263, Mar. 13, 2010. DOI: 10.1007/s10551-010-0463-8.

[43] H. P. Marvel, "Leegin and the economics of resale price maintenance," *Review of Industrial Organization*, vol. 50, no. 2, pp. 221–244, Feb. 7, 2017. DOI: 10.1007/s11151-017-9570-3.

[44] J. D. Iarrobino, "Turnover in the advancement profession," *International Journal of Educational Advancement*, vol. 6, no. 2, pp. 141–169, Jul. 14, 2006. DOI: 10.1057/palgrave.ijea.2150013.

[45] J.-P. Pirnay, E. Baudoux, O. Cornu, *et al.*, "Access to human tissues for research and product development," *EMBO reports*, vol. 16, no. 5, pp. 557–562, Apr. 7, 2015. DOI: 10.15252/embr.201540070.

[46] M. Á. González-Block, "Leadership, institution building and pay-back of health systems research in mexico.," *Health research policy and systems*, vol. 7, no. 1, pp. 22–22, Sep. 30, 2009. DOI: 10.1186/1478-4505-7-22.

[47] I. Hasan and M. Marinč, "Should competition policy in banking be amended during crises? lessons from the eu," *European Journal of Law and Economics*, vol. 42,

no. 2, pp. 295–324, May 12, 2013. DOI: 10.1007/s10657-013-9391-2.