

Declaration of Originality

I am the originator of this work and I have appropriately acknowledged all other original sources used as my references for this work.

I understand that Plagiarism is the act of taking and using the whole or any part of another person's work, including work generated by AI, and presenting it as my own.

I understand that Plagiarism is an academic offence.

and if I am found to have committed or abetted the offence of plagiarism in relation to this submitted work, disciplinary action will be enforced.

1. Documentation: Problem Framing (No coding required, PDF format, 25%, 50 marks)

1.1. Business Context & Decision (recommended: 1 page, 15 marks)

My generative-AI prompt:

“Draft me a Business Context & Decision for an interpretable salary benchmarker. Additionally, specify one decision it supports, who uses it, when it will be used, and how the final output will guide the actions. Ultimately, justify why this decision is prioritised while keeping it measurable and policy-oriented.”

The brief summary of AI's draft:

After the prompting, the AI proposed a quarterly publication of the salary bands and budget forecast and adjustments, which emphasises the accuracy, reliability, and transparency. It highlights market shifts as macro factors, and demand fluctuates; however, it lacks explicit ownership based on fairness and accuracy. Additionally, it provided a clear business plan and aftercare, such as approval and monitoring respectively. Ultimately, I kept the core idea but switched up the model's mechanics for deployment.

The key decision of the model will be made once per quarter and ad hoc if the market moves unexpectedly, which will be made on an impromptu basis. The organisation will decide whether to approve and publish the updated salary benchmark bands based on level, role, and skill cluster. Additionally, how to adjust training seats and budgeting towards the skills needed that command a higher salary.

Who will use it?	• Decision Owner: Head of the Labour Market, who reviews and signs off publication and the budget forecast/liquidity • Policy users: Skills planning and the budget office, who approve bands into funded seats and pricing of courses • Assurance: Ethics, governance, and communication, where they do fairness checks, disclosure to the public, and announcements
When will it be used?	• Cadence: Quarterly review in the second week of each new quarter • Early-warning triggers: Implement an interim hurdle when a sharp and unexpected pay shift occurs for any job role
How will it guide their actions?	• Evidence peak lands: Where each role/level/skill cell will have the updated band with the confidence range, conclusion/insights, data coverage, fairness checks, and stability/forecast over time • Budget mapping: Increase the funded seats and provider targets for the updated skill, showing the premium and demand

Fastest path from the model to policy

Many other analytics tools, such as provider performance or F1 scores, convert salary benchmarking to high-leverage actions. This action can serve as authoritative public guidance and inform targeted training investments. Ultimately, shaping employer expectations and behaviour for better choices and decision-making.

System-level signal correlation

The interpretable bands sharpen the price signals in the market, reducing under- and over-pricing and driving enrolment toward skill development with proper economic sentiment that does not require any regulatory intervention.

Governance and Accuracy

Since the decision is supported by explainable evidence and properly backed data, it's defensible to auditors, the public, and providers. Hence, this process visibly balances the accuracy, fairness, reliability, and stability.

Risks

- Benchmark guidance: Employers and future employees can anchor on outdated pay and salary growth progression
- Budget optimisation: Funds used in oversupplied, overvalued, low-return courses can be reduced, since high-value courses are capacity-constrained

In a nutshell, this decision is first because it is actionable and recurring. This model outputs to concrete levers while maintaining fairness and quality, based on responsiveness to market demands.

1.2 Task Definition: Target & Deployment Goal

The modelling task is to build an interpretable listing-time benchmark that estimates a fair pay for a job post based on role, level, skill cluster, month, work arrangement, and location.

Target:

- Primary: $\log(\text{salary_usd})$ at the listing, which stabilises variance, allowing robustness to outliers
- Secondary: an ordinal salary band, which is derived from the model's prediction, and is used for validation and not as a predictor

Target controls:

- Using only the listed salary fields and excluding the post-derived pay to avoid leakage
- Normalise the currency range and drop the irreconcilable records

Predictors for decision-making

Level, work arrangement, metro/region, posting month, market/talent indices, skill_count. Hence, these attributes are the main predictors that enable the model to provide the best prediction and benchmark for MOE/MOM.

Deployment goals for output triggers

Production of cell-level bands with intervals and driver cards for each quarter can trigger government decisions:

- Approval and publishing of eligible cells
- Rebalance training seats and budget toward skill with persistent premium and demand

Operational structure

Quarterly refresh with shock-triggered interims, which allows for continuous performance while maintaining fairness.

1.3 Formulation: Predictors & Feasibility

Three decision-critical predictors

Job level (Monotone driver)

- Decision value: Job level is a significant, policy-legible determiner of pay, where it anchors fair bands and prevents role inflation
- Has the strictly increasing effect, which is the monotonicity that enforces ordinal encoding as ordered categories. Hence, spotlighting the level steps in the driver cards
- Feasibility: column exists as job_level, which has low missingness and a stable taxonomy. Additionally, with low governance cost, it's a simple subgroup MAE parity checks by levels

Job skill + skill count (Monetisable edge)

- Decision value: Colonialised skills such as SQL, Power BI, Pytorch, and Python reveal the pay premia that directly lead to the training seats. Additionally, skill_count can distinguish a checkbox listing from a deep stack
- Parsing job_skills into 20-40 binary flags with the addition of skill_count, while rarer and uncommon skills will be placed into “Other”. Hence, the feature card shows the premia by bundle (Analyst + SQL + Python)
- Feasibility: dictionary-based parsing, where it's easy to keep current and track demand shifts

Work Arrangement (Pricing power based on supply/demand)

- Decision value: Pay varies with flexibility and competition, where it's crucial for publishable guidance
- Arrangement affects heterogeneity by tightening, such as remote/hybrid work arrangements appearing in hot markets but discounting in calm markets
- Combination of One-Hot for mode and supply/demand ratio from market_demand_idx and talent_supply_idx, which is a simple interaction in the model and driver stating the stringest premium where demand is hot
- Feasibility: All the relevant fields are present with low cardinality and are straightforward to monitor from quarter to quarter

Supporting controls:

- First_seen: Quarterly/Monthly for seasonality so bands are not biased by one-off spikes

Exclusions and justification:

- Salary_band: not a predictor as it's only used to validate published bands and data leakage risks
- Job_link: identifier that is unstable and needs to be dropped
- Company: high risk of single-firm overfit, which is based on size and sector, that may deviate from the average

1.4 Success Criteria and Environment

My generative-AI prompt:

“Draft me the Success Criteria & Environment for an interpretable, lisintg-time salary benchmarker. Hence, define the deployment KPIs, gates, monitoring rules, and list 3 real-world external factors with mitigation solutions, where all needs to be concise and measurable.”

The brief summary of AI's draft:

The AI suggested MAE tracking, with coverage and adoption, and flagged macros. I feel that it's useful, but too generic to use MAE, which has little to no fairness and stability, with no rollback policy when metrics are breached. Ultimately, I kept those ideas but added a north-star KPI with interim and advisory rules

North-star KPI – Seat Reallocation Effectiveness (SRE)

To see the deployment success, we see if the budget moves triggered by the model eventually pay off. SRE highlights the outcomes, which are the accuracy of seats based on the model, based on real market results. Thus, analysing the placements at or above the published band median or wage increment within 90 days to see the success of the model. Ultimately, aligning the model results based on policy which creates a simple feedback score for each bundle for easy audits and actions towards the funding.

- SRE is the percentage of the model-driven added seats achieving placement more than the band midpoint or wage increment based on 90 days
- Decision: If bundle SRE% % is based on the target, then further allocation will have an action

SRE Outcome	Decision	Action
SRE less as the target	Pause or reduce	Auto-pause the further allocations for the skill bundle and resume after only 1-2 clean cycles
SRE is same than the target	Maintain	Keep the current seats and monitor them quarterly while aiming for incremental improvements over time
SRE is more than the target	Increase	Scale the seats of the skill bundle if quality gates pass, like the MAE of wages

3 real-world external factors

Technology & Regulatory Injections

The adoption of AI and compliance changes can reprice and replace specific skills before historical data catches up. Therefore, employers can add new skill tags within a quarter

- New operations of a versioned skills dictionary (v1,v2,v3) and a process to add and remove skills
- Temporary limit reallocations for the new skills until stable placement and wage evidence is given
- Talk with major employers to validate the direction before commitment of major funds

Cross-border labour rules & FX

The changes of visa policies, remote work rules and regulations, and the fluctuation of the exchange rate of the USD/SGD pair, which alter the talent pools and salary compensation for roles such as remote and hybrid. Ultimately, observing the divergence of the premiums and salary shifts while tracking the FX rates can help employers revise their post-policy offers.

- Maintain mode-specific bands and add FX sensitivity note for accuracy and reliability
- Quick quarterly-based survey on remote offers for employers' pulse check

Macro events & Funding shocks

Based on the macroeconomic events, either stimulus bursts, unemployment rate fluctuates, hiring freezes, or FX rates change, all can push vacancies and wages faster than any quarterly cadence. Hence, sharp and sudden shifts in volume and offer median that do not align with the model's drivers will be cancelled.

- Operate an early-warning system that will see the recent volume, trends, and crucial financial news
- During high market volatility, issuing advisory guidance will help