Solutions

2.1

- The majority region dominates the pooled counts – the minority's mistakes are diluted to decimals ✓.

2.2

- In the data: the thinner minority class under-serves the boundary ✓.

2.3

- Equal-error-rate parity vs demographic parity (any two of the named criteria that conflict) ✓.

3.1

- Table: city 0.90, rural 0.75, mountain 0.50 – worst disparity 40 points (city vs mountain) ✓.
- Cause: mountain lates are $6/40 = 15\%$ of records – thinnest class; the boundary is fitted mostly on city data ✓.
- Data-side fix: collect more mountain records (or rebalance the sampling) so the class is not 15% ✓.
- Class weighting: weight mountain lates more in the loss – risk: precision in mountain may collapse (more false lates) ✓.
- Card: record the three-region recall table and the 40-point gap as a measured disparity ✓.

3.2

- It implements equal-error-rate (recall) parity by separated thresholds ✓.
- Right call: the cost of a missed late is the same per customer – parity matches the harm ✓.
- Mislead: thresholds tuned on small validation counts can overfit; and parity here can break precision parity, trading one harm for another ✓.
- Evaluation: per-region precision AND recall with confidence intervals on held-out data, before and after the move ✓.
- Duty: the limitation statement must disclose the separated thresholds and the measured residuals either way ✓.