Additional file 1. Complete model specification

The three primary MIMIC-IV logistic models and the two new dose-free models used for exploratory NWICU and INSPIRE external assessments were fitted in the same 2,840 MIMIC-IV patients (413 laboratory-recorded hypokalemia events). The entries below are **apparent fitted coefficients on the raw predictor scale**. Bootstrap optimism correction changed only the reported performance estimates, not these coefficient sets. The two dose-free models are separate fits; neither is obtained by deleting a dose term from the six-variable model.

Outcome, time point, and probability rule

The probability is for a valid potassium result <3.5 mmol/L sampled from 12 to <48 hours after the qualifying index record, before discharge or death in the index admission. These models were fitted among patients who were alive in ICU at 12 hours, had no known earlier low potassium, had an early potassium result for the model comparisons, and had enough later testing to classify outcome status. See the main Methods for the full cohort and laboratory rules. For each model, compute LP = intercept + Σ(coefficient × input) using Table S1, then $p = 1 / (1 + \exp(-LP))$.

Inputs and fitted coefficients

Potassium is entered in mmol/L, creatinine in μmol/L, and the recorded index furosemide dose in mg. To reproduce the fitted creatinine term, divide the SI creatinine value by 88.4 before taking log2; this recovers the numerical mg/dL input used for fitting. No other coefficient or intercept changes. The original age calculation is capped at 91 years before use. Recorded sex is coded male=1 and female=0. Creatinine and dose are transformed with base-2 logarithms **before** multiplication by the table coefficient. Dose is the positive amount recorded for the index event, not a cumulative delivered dose. "Baseline" refers to a sample taken before the index whose result had returned by the 12-hour landmark; it need not have been visible before the index. The early potassium input is the most recent post-index result returned by the landmark. The score does not include potassium change as an independent term.

Table S1. Intercepts and raw-scale coefficients for all fitted models

Term	Potassium only	Baseline clinical	12-hour updated	Dose-free baseline	Dose-free updated
Intercept	1.3175247651	-0.9451953090	2.8959009472	1.8439051526	5.5759009530
Baseline potassium (mmol/L)	-0.6912000262	-0.7060027359	-0.0638424930	-0.7901191719	-0.0707408804
Age, capped at 91 years	—	0.0024148237	0.0011817971	0.0025174406	0.0012535256
Male (1=yes, 0=no)	—	-0.4767301943	-0.4631436798	-0.5803551360	-0.5381872077
log2 baseline creatinine (μmol/L ÷ 88.4)	—	-0.0942604568	0.0619102870	0.2776937126	0.3785548306
log2 recorded index furosemide dose (mg)	—	0.4729329247	0.3938065499	—	—
Latest potassium result available by 12 hours (mmol/L)	—	—	-1.4726853742	—	-1.6253104033

The — symbol means that a term is absent from that model. Coefficients are provided to ten decimal places for reproducibility, not to imply clinical precision.

Missing values and implementation

In the MIMIC training sample, baseline creatinine was missing for two of 2,840 patients. The fitted training pipeline imputed **log2(creatinine)** to 0.1375035237, equivalent to approximately 97 μmol/L (1.1 mg/dL), before standardization; one of the 369 NWICU assessment patients, one of the 132 INSPIRE cardiovascular

patients, and eight of the 247 broader INSPIRE patients used this same fixed MIMIC median. No training index dose was missing. Baseline and early potassium, age, recorded sex, and a positive index dose where required were available under the analysis rules. The same MIMIC-fitted imputation and scaling were applied to both external datasets without refitting or recalibration. Raw-scale coefficients in Table S1 already incorporate the internal standardization, so a reader using the probability rule should **not standardize inputs again**.

The fitting pipeline used median imputation, standardization, and scikit-learn logistic regression (solver=lbfgs, C=1e6, max_iter=1000); fitting code is retained by the authors. The presented equations reproduce the model fits to the precision displayed. No clinical action threshold was developed. Figure 3 displays only the primary six-variable 12-hour-updated model, over approximately the development sample's 1st–99th percentile input ranges. Its visible total-point range maps to approximately 0.3%–69.1% predicted risk. Values outside those plotted ranges require the equation and were not assessed as a clinical calculator.

Scope of external assessment

The NWICU and INSPIRE comparisons used only the last two dose-free columns of Table S1, fixed after MIMIC fitting. NWICU Applied medication status, INJECTION prescription route, and laboratory storetime were used as recorded-data proxies; equivalence to MIMIC treatment and result-availability fields is unconfirmed. INSPIRE records IV administration, but its laboratory time is a sampling-time proxy and its values are binned for deidentification. The original six-variable model and Figure 3 were **not externally validated** in either dataset. Neither the equations nor the nomogram should be used to direct care without prospective evaluation and reliable field harmonization.