Abstract. We study imputation in large dependent panels when missing observations are resolved sequentially and the completed panel is subsequently used for multivariate forecasting. We develop contextual bandit imputation, a data-adaptive rule that selects among univariate, factor-based, and state-space imputers using only information available at the decision date and reveal-based feedback. The framework accommodates serial and cross-sectional dependence, batched updating, and contextual heterogeneity, and delivers regret bounds relative to a contextual oracle. In Monte Carlo designs calibrated to ragged-edge panels, the procedure improves the trade-off between cellwise imputation accuracy and forecast performance. In a real-time FRED-MD application, it matches PCA locally, improves upon it in nowcasting, and remains close to the strongest univariate benchmark.