SONNETS : Medium-Term Problem Definition Document

EPSRC Programme (EP/X036006/1)

David Thomas, Wayne Luk, Rishad Shafik, Andrew Brown
Mark Vousden, Graeme Bragg, Alex Yakovlev, Kabita Adhikari, Walter Distaso
Shidur Rahman, Alex Chan, Tousif Rahman, Ce Geo, Paul Labonne, Bob Pattison
18th Aug 2025

Introduction

The purpose of this document is to provide a medium-term problem definition for the SONNETS project, in terms of specific research objectives and delivered capabilities over the next year to the next advisory board meeting in May 2026. This document is both part of the original SONNETS work-plan, and was also requested by the SONNETS advisory board after first year review.

The document does not attempt to define research beyond one year ahead, as there is currently significant ongoing engagement with financial end-users and stakeholders to explore the most relevant and impactful research topics in terms of financial risk analysis.

This document does not describe all research within the SONNETS project, and focuses only on the integrated delivered system that demonstrates work from across the project in the context of national financial risk analysis. In the rest of the document this is referred to as the "integrated system".

High-level financial risk analysis objectives

From a financial end-user perspective, the goal of the integrated system is to provide an automatically updated analysis of real-time UK financial risk, including types of analysis and risk estimates that are not currently available. In this context "real-time" means automatic daily updates of risk analysis results, and is intended to contrast with the much slower weekly, monthly or manually-driven updates commonly used in academic and regulatory risk analysis systems.

The notion of financial risk is deliberately broad and includes:

- Quantitative forecasting of key financial variables including macro-economic factors (GDP, inflation, and interest rates) and financial market indicators (stock market indices, volatility indices, and bond yields).
- Identification of causal relationships between key financial variables and also with other financial input variables, such as the impact of interest rate changes on inflation or GDP, or the impact of stock market volatility on bond yields.
- Generation of plausible and worst-case scenarios for changes in input variables based on historical data and machine learning models, and automatic evaluation of the impact of these scenarios on key financial variables.

The technical requirements and features of the integrated system showcase novel capabilities that are not currently available in combined financial risk analysis systems. Particular capabilities we will showcase are:

- A data-driven approach that identifies potential causal relationships between financial variables, rather than relying on human experts to construct or estimate these relationships.
- Scenario generation based on machine learning models, which use historical data and learnt models to generate a selection of plausible and worst-case scenarios.
- Evaluation of the impact of these scenarios on key financial variables using SVAR (Structural Vector Autoregression) and causal graphs, without requiring extensive simulation or human intervention.
- Estimates of risk using novel machine-learning models, which provide both high quality predictions and estimates of uncertainty, as well as explanations of how the predictions were made.
- An automated always-on approach that continuously updates the analysis and risk estimates, rather than requiring manual intervention or periodic updates.
- A single integrated system that combines both model learning and inference/prediction based on those models, rather than using separate systems for offline or batch learning.
- Ingestion of additional non-quantitative data sources such as news articles and financial statements to provide additional inputs to both causal graph analysis and scenario generation.
- A web-based dashboard that provides real-time updates and visualisations of the risk analysis results, making it accessible to a wide range of users, including a public view for outreach purposes.

Optional features that may be included are:

- Interactive what-if modelling, allowing users to choose specific scenarios and see the impact on key financial variables recalculated in real-time based on the latest SVAR and causal graph analysis models.
- Intra-day updates of the risk analysis results based on the arrival of new data, allowing users to see the latest risk estimates and causal relationships at an hourly level.
- Integration of hardware accelerated algorithms to speed up the analysis and learning processes, such as using FPGAs or GPUs for causal graph analysis or Tsetlin Machine model training.

Integrated System Requirements

From an end-user point of view, the integrated system appears as an always-on, continually updated web dashboard integrating a variety of novel and classical risk analyses methods updated and recalculated at a daily level. This system will be deployed across a combination of commercial and private cloud infrastructure, and automatically scale up or out according to analysis demands (e.g. if market conditions change substantially), but will scale down most of the time for cost and power efficiency.

The specific requirements for the integrated system are mostly functional (i.e. what it should do and the capabilities it should provide), with non-functional requirements imposed in terms of update frequency and system up-time. Other non-functional aspects such as compute efficiency, ability to handle higher data volumes are important and will be measured as part of the research, but no specific requirements are imposed at this stage. The requirements are intended to be well-defined enough to definitively say

whether and when they have been met, but also sufficiently open-ended to allow for exploration of novel approaches and techniques.

The integrated system is intended to be delivered as two milestones:

- Milestone 1 - due Dec 2025 : Integrates existing SONNETS research components and capabilities developed in the first year of the project.
- Milestone 2 - due May 2026 : Integrates new SONNETS research components and capabilities currently being developed in the second year of the project.

The phase 1 integration is already in progress, and multiple components needed for phase 2 are already being explored as research topics.

The overall integrated system is shown below, with components and functions for milestone 1 shown at the top half, and those for milestone 2 in the bottom. The diagram shows only the coarse information flow within the system and not attempt to show all the individual data flows between components. Components in red and blue represent existing research completed during year 1 of SONNETS, while purple flows are combined ETAI+ETM research currently being developed. Research from ETC is not shown, but is implicitly supporting the design, integration, and deployment of the overall system.

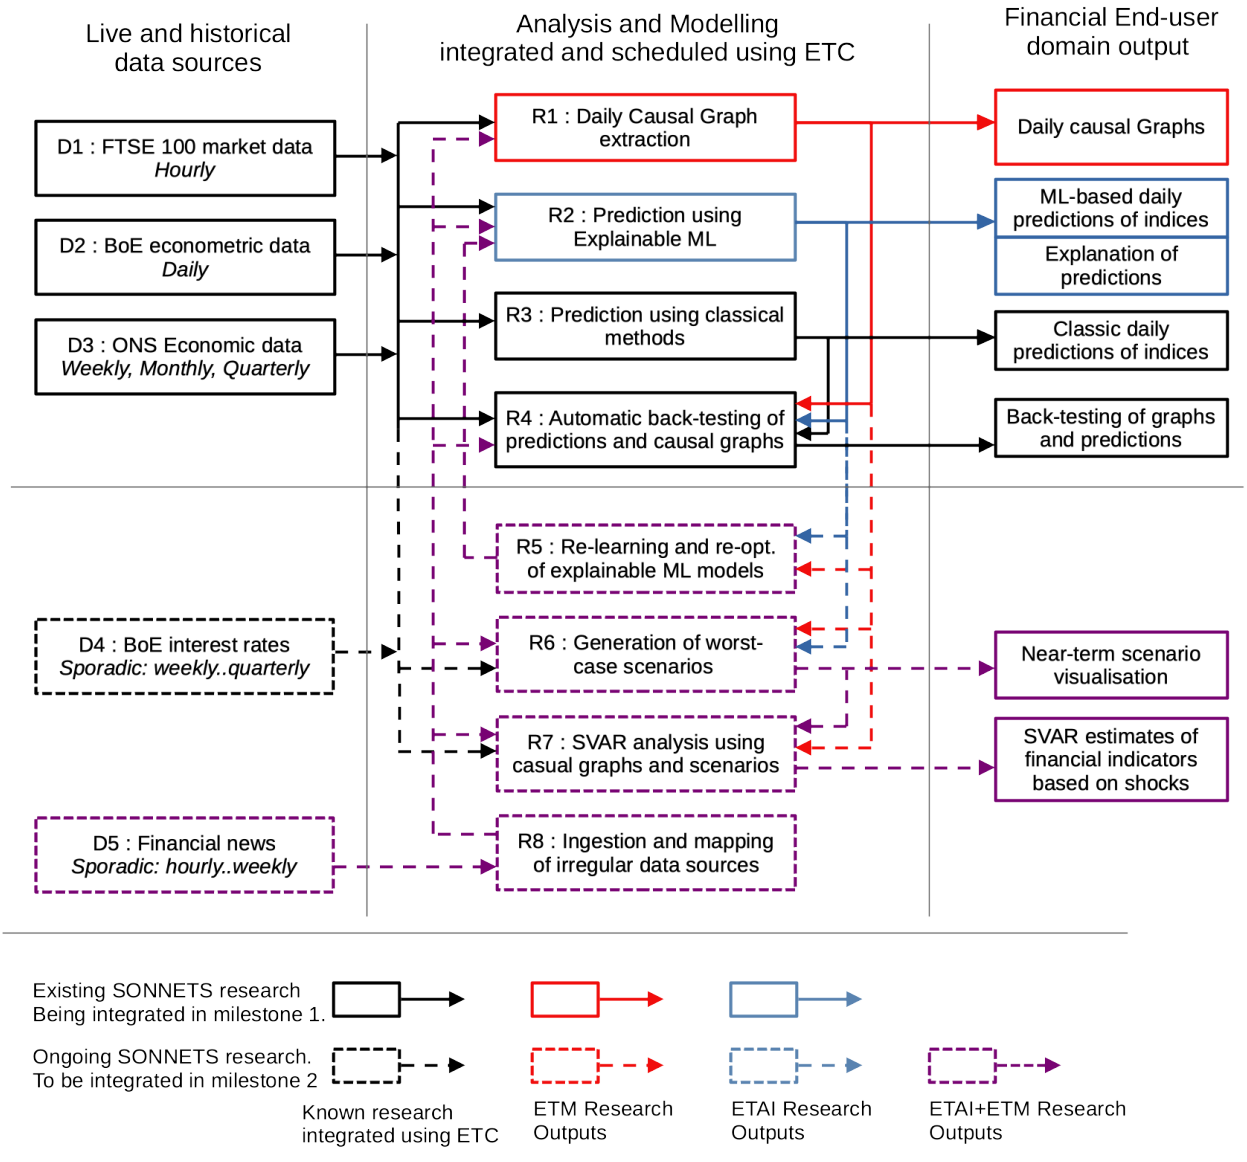

General system requirements

Milestone 1:

- G1 : Automated ingestion of data from live and historical data sources. *Ability to work with real-time data.*
- G2 : Live casual graph and risk analysis results to be updated at least daily. *Evidence for lack of manual intervention.*
- G3 : Display of past risk analysis results. *Used to back-test the system and evaluate quality of results.*
- G4 : Achieved system uptime of at least 7 days. *Initial proof of robustness and reliability.*

Milestone 2:

- G5 : Autonomous on-demand re-calculation of analysis and modelling functions. *Ability to integrate both short-term analysis and long-term learning in one system.*
- G6 : Display of the latest risk analysis results on a publicly accessible website. *For outreach and engagement purposes, and as evidence system runs autonomously.*
- G7 : System deployed and operating within standard cloud infrastructure. *Show that the combined multi-analysis application is deployable in practice.*
- G8 : Achieved system uptime of at least 31 days for all components from milestone 1. *Demonstrates ability for system to run autonomously.*

Risk analysis requirements

Milestone 1:

- R1 : Daily estimation of causal-graphs for key financial indicators across multiple time-series (D1..D3). *Demonstrates scalability of causal graph analysis.*
- R2 : Daily predictions of key economic variables for one month ahead using explainable machine learning applied to time-series (D1..D3) based on STRIPE-based models and Tsetlin Machines. *Demonstrates ability to create explanations along-side models.*
- R3 : Daily predictions of key economic variables for one month ahead using standard econometric models (e.g. ARIMA, GARCH, etc.) and machine learning (e.g. xgBoost, random forests) from existing libraries. *Demonstrates ability to include multiple analyses including in-house and external libraries in a unified application.*
- R4 : Automated back-testing of estimates and predictions from R1, R2, and R3 to automatically assess quality of results (feeding G5). *Provides quantitative metrics for quality of results and comparison with existing systems.*

Milestone 2:

- R5 : Daily incremental re-learning and re-optimisation of explainable machine learning models used by R2. *Demonstrates ability to automatically update models without manual intervention.*
- R6 : Scenario generation of plausible and worst-case scenarios for key economic variables, based on historical data and explainable machine learning models. *Demonstrates ability to generate scenarios*

without human intervention.

- R7 : SVAR analysis of key economic variables using the causal graphs from R1 and the scenarios from R6. *Demonstrates ability to perform composite analysis incorporating both AI and modelling.*
- R8 : Addition of new data sources (D4,D5), such as infrequent but high-impact interest rate changes and non-time-series news articles. *Demonstrates ability to integrate data sources with very different frequencies, modalities, and impact.*

Data sources

In the near term we focus on data sources that are publicly available:

Milestone 1:

- D1 : FTSE 100 hourly data.
- D2 : Daily, weekly, and monthly data from the Bank of England (via bankstats).
- D3 : ONS macroeconomic data (e.g. GDP, inflation, interest rates, etc.) at multiple time scales.

Milestone 2:

- D4 : Bank of England interest rate changes (very infrequent).
- D5 : News sources describing major economic events.

Due to cost and availability of live data, the live demonstrator may use historical data from one month ago, but played back in real-time. Historical data will be used for back-testing and offline evaluation of the demonstrator system.

Conclusion

The requirements outlined in this document will result in a system that incorporates a wide range of novel and classical risk analysis techniques, and provides a real-time, always-on analysis of UK financial risk. As an integrated system this provides capabilities that do not currently exist in an academic platform for national financial risk analysis, and forms the basis for even more advanced risk analysis in the future.