Improving Econometric Forecasting: Functional Analytic Fixed Point Methods for Developing Hybridized Structural Models

Eliot Ozaki

6/5/2026

Abstract

This paper is an exposition on the challenges of modern econometric modeling, and how econometric models either are too rigid or too unexplainable (Iskhakov, 2020; Kilic, 2025). We see this distinction much in the difference between traditional structural econometric models and machine learning models, where the traditional models often struggle for accuracy whereas the machine learning models struggle to output any real causal information (Iskhakov, 2020; Woloszyn & Bukowski, 2025). We develop a strong argument towards the use of hybrid models for econometric estimation rather than either option alone, models which address the weaknesses of both without developing any new major pitfalls (Buckmann et al., 2021; Lee, 2025). Within this paper, we also outline the necessary tools for the development of a type of model based on posing economic questions as inverse problems, which is well explored in Carrasco et al. (2007). This outline includes a summary of the inverse problems framework of Carrasco et al. (2007), an explanation of the Morozov Discrepancy Principle (Engl et al., 1996; Morozov, 1966, 1984) which is one of the key steps in the model, and finally an overview of our original model algorithm for recovering latent economic variables through estimation with Landweber regularization (Engl et al., 1996; Hanke et al., 1995; Landweber, 1951).

1 Introduction

Economics, as a subject, is rooted in theory, from the most basic research to ground-breaking policy development. Early in university, students learn the basics of supply and demand alongside fundamental concepts like utility maximization, equilibrium, and rational expectations. These ideas are not just classroom abstractions, but the backbone of the structural theoretical models that guide monetary and fiscal policy.

Traditional, theory-based economics has worked for many decades and has proven itself in both practical settings and academic research. However, one weakness has repeatedly emerged in this traditional, structural version of economic modeling: when the economy or key variables do not align with rational trends or steady-state assumptions, theoretical models fail. There are clear historical examples of this pattern. After World War I in the United States, during the 2007 housing crisis, the 2012 technology bubble, and in many other instances, the models that guided policy either missed crucial turning points or had to be replaced or heavily revised, as with the shift toward Keynesian and Hayekian perspectives in the post-war era (Kilic, 2025; Pricepedia, 2025). This inability to adapt to irrational, high-pressure, or otherwise contentious economic contexts reveals a serious weakness in purely structural modeling, especially when it is used to design policies aimed at mitigating or avoiding the worst effects of shocks or crises.

Over the last two decades, a more modern methodology of machine-learning-based econometrics and data-driven economics has emerged, challenging the dominance of traditional models. Machine learning (ML) has developed to the point where, in many forecasting tasks, it outperforms classical econometric approaches, particularly in detecting complex, nonlinear patterns in large datasets. In economics, this has opened up a new set of tools for understanding and forecasting macroeconomic and financial variables (Batool, 2022; Kristjanpoller, 2024).

Yet ML approaches have their own limitations. Economists value explainability and theoretical coherence, yet many ML models lack them. There is also skepticism about whether highly flexible models generalize well or simply overfit noise, especially in unstable environments.

This raises a central and urgent question; is it possible to implement machine learning techniques inside traditional structural economic models in a way that increases robustness and interpretability without sacrificing flexibility and forecasting accuracy? Can we design hybrid models that retain the credibility of theory while leveraging the adaptability of ML? This paper argues that hybridizing structural econometric models with machine learning can make macroeconomic forecasts more credible and statistically robust. To address this question, the paper first examines the limitations of both structural and pure ML models, then defines what constitutes a genuine hybrid model, and finally evaluates the effectiveness of such hybrid approaches in enhancing forecast reliability, interpretability, and adaptability. By integrating ML's flexible pattern recognition into well-understood theoretical frameworks, hybrid models aim to deliver reliable, interpretable, and responsive forecasts. These models directly address the tension between interpretability and adaptability in policymaking, so it is imperative that researchers and practitioners actively pursue, implement, and refine such hybrid approaches to drive more responsible and effective economic forecasting for the future.

2 The Problem: Inadequacies of Current Methods

2.1 Structural Models: Strengths and Failures

The fact that the economic models used at the highest levels of government may require further development is contentious, so it is necessary to explain clearly why this is the case. In general, traditional structural econometric models emphasize credibility and interpretability. They are carefully designed based on well-established economic theory to capture macroeconomic or microeconomic signals and generate outputs that inform policymakers, central bankers, and researchers about markets and the broader economy. This theoretical foundation gives structural models significant weight and perceived robustness, and, because of their intentional structure, they provide economists with considerable explanatory power (Iskhakov, 2020; Woloszyn & Bukowski, 2025).

Yet like most models, structural systems face a tradeoff. There is an implicit tension between explainability and accuracy. These models rarely capture the full motion of the economy. In practice, finding patterns that perfectly match observed dynamics is unusual for a general structural model. They often lack the specificity or complexity needed to account for the entire economy and its heterogeneous agents. As a result, economists often work with models that do not fully capture the market economy, simply because they are built on simplifying assumptions that limit accuracy in practice (Kilic, 2025; Pricepedia, 2025).

Certain assumptions used in these models may also be inaccurate or non-generalizable. To have a structural model, one must assume a structure, which often implies linear functional forms or stable parameters. Yet many relationships that are assumed to be linear in theory turn out to be nonlinear in reality. When non-linear relationships are represented with linear functions in estimation, this mis-specification naturally produces errors and inaccuracies. Furthermore, structural models sometimes encode incorrect or overly strong assumptions about causality and correlation. For instance, in difference-in-differences (DiD) econometrics, a treatment variable is sometimes assumed to affect an outcome directly, when in reality the relationship may be indirect or driven by shared shocks. Misinterpreting correlation as causation leads to biased or misleading structural results (Batool, 2022; Kilic, 2025).

The main problem of structural models can be summarized in a single word: structure. They attempt to measure or read a dynamic, largely non-structural entity, such as the real economy, using a framework that requires rigidity and pre-specified relationships. This rigidity results in forecast lags, inaccuracies, and an inability to track economic behavior in high-stress situations, such as abrupt shocks, crises, or regime changes. Structural models remain strong in their explainability and in contexts where assumptions hold, but there is a clear issue with their use for forecasting or extrapolation into an uncertain future (Batool, 2022; Iskhakov, 2020; Kilic, 2025).

2.2 Pure Machine Learning Models: Power and Opacity

Alongside structural models, there has been a significant increase in the use of strictly data-driven, machine-learning-based models in economics. These approaches almost invert the pattern of strengths and weaknesses seen in structural models. Pure ML models have little to no built-in economic theory and, as a result, often lack interpretability in classical terms. At the same time, they can be much more accurate than traditional models in a wide range of settings, including crises, because they are designed to detect patterns in the data rather than enforce predetermined relationships (Buckmann et al., 2021; Iskhakov, 2020; Woloszyn & Bukowski, 2025).

In practice, such models are already used by institutions such as central banks and financial regulators to support forecasting and risk assessment. Their advantage is the ability to provide an accurate read on how the economy or a particular market is moving, often in near real time. This gives policymakers access to improved short-term projections, which are especially useful when rapid action is required, for example, in financial stress episodes or sudden commodity price shocks (Batool, 2022; Gietner, 2025; Lee, 2025).

However, this strength leads directly to the core concern about ML models: explainability. It is tempting, especially for inexperienced analysts, to assume that ML models cannot be wrong because they use large amounts of "hard data." Without proper skepticism, they may be applied uncritically. The black-box nature of many ML techniques means they can pick up and exploit correlations or patterns without regard for whether these patterns have economic meaning or causal interpretation. Theory may indicate that no relationship exists between two variables, even when the data show a strong correlation. ML models do not know this; in many cases, they will simply use such spurious relationships as if they were genuine signals (Buckmann et al., 2021; Hammann & Wouters, 2025).

This disconnect between predictive performance and economic logic is a growing concern in the use of ML in economics, especially when models inform policy. A forecast that cannot be explained or defended in terms of theory and evidence is unlikely to be trusted, no matter how accurate it appears to be in backtests. Lack of explainability thus becomes a barrier to adoption in public institutions and undermines accountability (Buckmann et al., 2021; Lee, 2025; Woloszyn & Bukowski, 2025).

2.3 The Gap Between Approaches

Taken together, structural and ML models offer complementary strengths and weaknesses. Structural models provide interpretability and credibility but can be inaccurate or slow to adapt in complex, changing environments. ML models provide powerful pattern recognition and adaptability, but are opaque and often hard to justify in policy-making contexts. Faced with this tradeoff, many economists and policymakers continue to default to structural models, as they are more trustworthy in the sense that their logic can be traced and defended (Iskhakov, 2020; Kilic, 2025; Woloszyn & Bukowski, 2025).

This does not mean that accuracy must always be sacrificed in favor of interpretability. Instead, the trade-off between these two types of models suggests a gap that can be filled by carefully designed hybrid econometric models. Such hybrids would aim to combine the strengths of each family while mitigating their most problematic weaknesses. The question is whether this is practically and theoretically feasible.

2.4 Broader Implications: Credibility, Policy Risk, and Structural Gaps

To move forward with the idea of hybrid models, it is useful to state clearly what is required for a forecasting framework to be considered successful in policy-relevant macroeconomics. Three central requirements emerge: accuracy, credibility, and interpretability. Other desirable properties, such as debiasing, ease of use, or computational efficiency, tend to fall under one of these headings or are secondary.

Accuracy is the most straightforward requirement. For research in this area to have tangible value, a new model or framework must match or exceed the performance of existing models in forecasting key variables, especially in environments where current methods perform worst. If a hybrid approach cannot improve accuracy relative to traditional structural or pure ML models, particularly during crises or high-volatility episodes, then its complexity may not be justified. In such a case, it might be more reasonable simply to continue using well-understood structural models (Batool, 2022; Kristjanpoller, 2024; Lee, 2025).

Credibility is just as important as accuracy, especially in policy settings. If only accuracy mattered, institutions could adopt the most accurate ML models, regardless of whether they were transparent. In practice, a prediction that cannot be trusted or explained is almost as useless as a prediction that is simply wrong. Credibility includes the perceived legitimacy of the model and the institution that uses it. It is necessary to maintain public trust and ensure that policy decisions are grounded in evidence rather than ideology. When a model becomes less credible or less clear about its legitimacy, its outputs lose viability in public and political discourse, and the resulting policies may be viewed as speculative or partisan (Buckmann et al., 2021; Pricepedia, 2025; Woloszyn & Bukowski, 2025).

Interpretability is closely related to credibility, but not identical to it. In government and central banking, the ease with which results can be communicated is critical. The most accurate forecast in the world has limited value if policymakers, staff, and the public cannot understand it. Models need to produce clear and interpretable results that allow economists to explain what is predicted and why. Without a convincing "why," it is difficult to place genuine trust in model-based results, even if they have been accurate in the past. Opaque economic predictions complicate parliamentary and central bank negotiations and can alienate non-technical stakeholders. At a more basic level, democratic accountability requires that citizens can, at some level, understand what their government is doing and why (Buckmann et al., 2021; Gietner, 2025; Iskhakov, 2020).

At present, there is no single modeling framework in common use that fully satisfies all three requirements. Structural models generally score well on credibility and interpretability but can fail on accuracy in volatile or nonlinear environments. Pure ML models can excel in accuracy but fall short in credibility and interpretability. This mismatch is the deeper reason that hybrid approaches are worth considering. The aim is to construct forecasting systems that are accurate, credible, and interpretable enough to be used in real-world policy settings.

2.5 Visualizing the Explainability–Accuracy Tradeoff

One of the main issues in econometric and statistical modeling is choosing the appropriate model for a given task. Common models such as linear regression, multiple linear regression (MLR), decision trees, and more complex machine learning methods, such as deep neural networks, differ along two key dimensions: explainability and accuracy. Explainable models produce results that are easy to interpret, enabling researchers and policymakers to clearly see which variables matter and how they are related. Accurate models are those whose predictions more closely match real-world outcomes. In practice, there is usually a negative relationship between these two features. Models that gain accuracy tend to lose interpretability, while highly interpretable models often sacrifice predictive performance. This tension lies at the heart of the research question: whether hybrid modeling approaches can help address the limitations of current methods in both routine and crisis conditions.

Figure 2.5 presents the tradeoff between explainability and accuracy for several representative forecasting models, adapted from Hammann and Wouters (2025). In the figure, models are placed according to their typical position along the two dimensions. Simple models such as linear regression and decision trees tend to perform well in terms of explainability but only achieve moderate accuracy. Their structure and parameters can be directly interpreted, and policymakers can see how changes in inputs affect outcomes. In contrast, models such as gradient-boosted trees and deep neural networks lie in the high-accuracy but low-explainability region. Their ability to capture complex non-linear patterns in data allow them to achieve clear superiority in predictive accuracy, but it comes at the cost of those complex patterns being uninterpretable in terms of economic causality or mechanism. The overall pattern shows that as models become more flexible and powerful in prediction, they tend to become less transparent and harder to justify in policy discussions (Hammann & Wouters, 2025).

This visual is important because it makes the central tradeoff that motivates the rest of the paper concrete. It provides evidence that there is no widely used, off-the-shelf

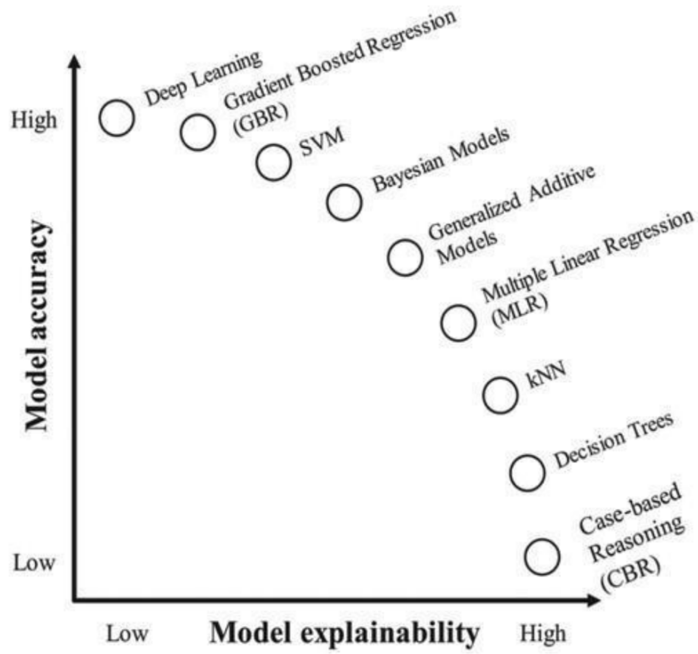

Figure 1. Explainability versus Accuracy in Forecasting Models
Source: adapted from Hammann and Wouters (2025).

model that is both highly accurate and highly interpretable. This is exactly the problem policymakers face when choosing a forecasting approach.

This figure supports the claim that neither extreme is sufficient on its own. Models that sit entirely on the explainable side risk being too simplistic or fragile in volatile conditions, while models that sit entirely on the accurate side risk being unusable in institutional settings that require transparency and justification.

The implication is that new approaches, including hybrid models and explainable AI methods, are needed to enter the space where both goals can be pursued simultaneously. The figure does not, by itself, prove that any specific hybrid design is superior to a particular middle-ground model such as MLR, but it clearly illustrates why balancing accuracy and interpretability is necessary, and why the research focus turns toward combining strengths from different modeling traditions rather than choosing a single method.

3 Defining Hybrid Econometric–Machine Learning Models

3.1 Building a Definition

To make progress on hybrid modeling, it is important to define the concept clearly. In econometrics, as in many academic fields, there has been a long-standing distinction between theoretical studies and real-world, data-driven research. Theoretical econometricians focus on building abstract models based on economic theory, often using simplified assumptions and simulated or highly controlled data to see how an economy should behave under those assumptions (Iskhakov, 2020). The goal of this work is internal consistency and clear logical structure, rather than perfectly matching real-world fluctuations. Applied econometrics, by contrast, focuses on models estimated primarily using real-world data with minimal implicit assumptions, emphasizing reflection of actual economic behavior. Applied economists and econometricians are concerned with how well a model fits or predicts observed outcomes and will often sacrifice some theoretical neatness to better capture complex relationships in the data.

As computing power and data availability have grown, it has become more realistic to combine a structural understanding of how the economy works with flexible, data-driven methods that learn from observed outcomes. Recently, a third group of economists and a third partition of econometric work has been developing in this direction: hybrid econometrics. This approach seeks to combine the interpretability and credibility of structural theoretical models with the accuracy and adaptability of models built directly from real-world data.

Formally, a hybrid econometric–machine learning model is a forecasting framework that systematically combines a structural econometric model with a machine learning model in a functionally interdependent manner. Two main features are required. First, the hybrid must include both a theoretical, structural component and a data-driven ML component. Second, the outputs, parameters, or structures of these components must directly affect one another in producing the final forecast. The integration must be explicit and meaningful. Hybrid models are not simply side-by-side comparisons in which econometric and ML models are estimated separately, and the best one is chosen. They are also not mere preprocessing pipelines in which ML is used to clean or compress the data and then discarded.

3.2 Examples and Boundary Cases

An ARIMA–LSTM model provides a clear example of a hybrid under this definition. ARIMA (AutoRegressive Integrated Moving Average) is a classical econometric time-series model that captures linear structure in the data. LSTM (Long Short-Term Memory) is a recurrent neural network type of ML model that can learn nonlinear temporal patterns. In a hybrid ARIMA–LSTM design, the researcher fits an ARIMA model to the data and uses its residuals as inputs to an LSTM. The ARIMA component handles predictable linear dynamics, while the LSTM learns patterns that the linear model cannot capture. The final forecast is the sum of the ARIMA prediction and the LSTM-based correction. Here, the models work together in a single pipeline, and the LSTM's behavior is driven by the econometric model's errors (Batool, 2022; Kristjanpoller, 2024).

In contrast, consider a study that runs a VAR and a random forest separately on the same macroeconomic dataset and compares which performs better. VAR is a standard econometric model that captures linear interdependencies, whereas random forests are flexible, nonlinear ML models built from ensembles of decision trees. Although both types of models are present, there is no interdependence: they do not inform each other or combine outputs. This is a comparison, not a hybrid.

Some cases might look hybrid but are not. For example, using PCA or clustering to preprocess data before running a structural regression involves ML methods, but they are not part of the forecasting mechanism itself. The regression model could be estimated in the same way regardless of the specific ML tool used for preprocessing. The ML component here is an assistant, not an integral part of the model.

Finally, hybrid behavior can occur within familiar frameworks. A two-sector DSGE model in which one sector's shocks are estimated using traditional Bayesian methods while the other sector's parameters are estimated by a neural network trained on microdata creates a tightly coupled system. The neural network's outputs feed into the DSGE's equilibrium and influence dynamic responses. Even though the high-level structure remains that of a DSGE model, the embedded ML block makes the system a genuine hybrid (Iskhakov, 2020; Smets & Wouters, 2007).

4 Solutions: Hybrid Modeling Implementation and Advantages

Building on this definition, the paper now turns to how hybrid models might actually be implemented in macroeconomic forecasting and why they can be advantageous. One useful way to think about structural models is as puzzles made up of smaller pieces. A DSGE model, for example, links together submodels for households, firms, fiscal policy, and monetary policy, each with its own behavioral equations and parameters. In practice, some of these pieces perform well during crises and structural breaks, while others fail. ML can be deployed strategically to support or replace the weaker parts, while preserving the overall theoretical structure.

In a DSGE or multi-sector setting, this might mean using neural networks or other ML tools to estimate shock processes or behavioral parameters in specific sectors that are difficult to model structurally, such as a financial sector with nonlinear leverage constraints, while keeping more standard structural representations for households and firms (Smets & Wouters, 2007). ML can also be used to estimate time-varying parameters or to classify regimes in which different linearized approximations hold.

Time-series forecasting offers another useful implementation path. ARIMA–LSTM hybrids demonstrate how classical linear time-series structure can be combined with nonlinear sequence modeling to improve forecasts of financial or macroeconomic variables (Batool, 2022;

Kristjanpoller, 2024). In these setups, the ARIMA model provides a transparent, familiar baseline for economists, while the LSTM corrects for systematic residual patterns.

Across the literature, such hybrid setups have delivered measurable improvements in forecast accuracy and often maintain a reasonable degree of interpretability when designed carefully (Buckmann et al., 2021; Lee, 2025; Woloszyn & Bukowski, 2025). Importantly, they tend to perform especially well during periods of instability, exactly when traditional models struggle and when policymakers most need reliable information (Batool, 2022; Lee, 2025; Pricepedia, 2025). At the same time, recent advances in explainable AI, such as feature importance analysis, scenario-based explanations, and local approximation methods, provide tools for rendering ML components more transparent than before (Buckmann et al., 2021).

In light of these developments, the proposal is to construct hybrid forecasting models with explicit modular designs that combine the strengths of traditional structural economics with the flexibility and performance of modern data-driven methods. If successful, such models could enable a new era of credible, interpretable, and accurate forecasting in economic policy and practice.

5 Hybrid Fixed Point Modeling

We thus turn to the mathematical portion of this paper, the actual construction of a hybrid model. Given the requirements of a flexible, accurate model, it makes sense to construct a model based on fixed point estimation. The key reason for this is the robustness of fixed point theorems; under key assumptions, fixed point theorems allow us to be certain of convergence to a limit no matter how poor our conditions or guess is, and given proper construction can give us incredibly accurate answers no matter the context.

The model described in this paper is based on the fact that structural econometric problems can, in many cases, be reconstructed as inverse problems. We follow closely with the language and ideas of Carrasco et al. (2007), and develop a model that starts with an initial operator, dataset, and unknown true latent state economic variable. We then use Landweber iteration to regularize our estimation, and machine learning to debias and ensure accuracy and credibility in our output. As a result, we not only construct a debiased, accurate estimate of the true latent state of the economy, but we also get information about how that latent state affects nonstandard economic variables.

One note to be said is that this model is heavily based on pre-existing structural models acting as a linear operator, meaning that most of the explainability of this model comes from that relationship rather than any new parameter or coefficient. The most we truly calculate in terms of coefficients or relationship parameters comes in the form of our machine learning operator, which can be manipulated to glean information on relationships between the economic variable and other relevant variables. We will first outline necessary

We will first introduce background around Hilbert spaces, the type of spaces the model and model parameters live in. We then review the Carrasco et al. (2007) literature as it is very relevant background to have when trying to understand out model. Finally, we will explain the methodology and clarify a specific method used within the methodology.

6 Preliminary on Hilbert Spaces

Prior to any calculation, it is important to define the relevant space we are working in and why we choose to work in this space. Given the economic context and the shape of the models and data, a Hilbert space seems to be the correct metric space to address our question in. To explain the specific reasons for this, we must first outline what a Hilbert space actually is, and the properties it holds. We begin with an outline of metric spaces as a whole, introducing relevant topics until the requisite features for a Hilbert space appear, allowing us to give an exact definition of the space and its benefits.

Beginning with metric spaces, a metric space is simply a set that is paired or equipped with a distance function (Axler, 2020). There is no restriction on the items of this set, and the distance function must only satisfy the following conditions. Given a metric space $X = (K, d)$, for all $x, y, z \in K$:

$$d(x, y) = 0 \iff x = y, \tag{6.1}$$

$$d(x, x) = 0, \tag{6.2}$$

$$d(x, z) \leq d(x, y) + d(y, z), \tag{6.3}$$

$$d(x, y) = d(y, x). \tag{6.4}$$

A basic but more extreme example of a metric space is $K = \{E, O\}$ with $d(x, y) = 0$ if $x = y$ and $d(x, y) = 1$ otherwise (Axler, 2020). We see that K contains just the letters E, O which have no mathematical meaning and only represent two elements of the set. Using two meaningless letters represents the lack of necessity for elements to be able to relate to each other. One could also substitute E, O for shapes like square and circle and retain the same property. Trivially, $d(x, y)$ has the conditions required.

Now define a vector space as a set V equipped with addition and scalar multiplication satisfying commutativity, associativity, additive identity and inverse, multiplicative identity, and the distributive property (Axler, 2020). Should any of these not hold, a space would not be considered a vector space. Further, we define a normed linear space as a vector space where the distance function is given by a norm, inducing a metric on the space (Axler, 2020). For instance, in $\mathbb{R}^n$ with the Euclidean norm $\|x\| = \sqrt{\sum_{i=1}^{n} x_i^2}$, we can meaningfully quantify the distance between two vectors and discuss convergence rates and error bounds.

Within normed linear spaces we now define two distinct spaces. Banach spaces are normed linear spaces with the restriction that all Cauchy sequences must converge, i.e., the underlying set is complete (Axler, 2020). This completeness imposes structure on sequences of points, allowing researchers to use convergence theorems they otherwise could not. We also define inner product spaces as normed linear spaces where the norm is defined by the inner product $\langle x, y \rangle$ (Axler, 2020).

Finally, we have the foundation to define a Hilbert space. A Hilbert space is a complete inner product space, equivalently a Banach space whose norm is defined by the inner product (Axler, 2020), with norm $\|x\| = \sqrt{\langle x, x \rangle}$ and distance $d(x, y) = \sqrt{\langle x - y, x - y \rangle}$. The inner product is specifically important because it allows us to use theorems like Cauchy–Schwarz (Axler, 2020) to dictate more about convergent sequences and find upper bounds for distances. In total, Hilbert spaces are complete vector spaces equipped with an inner product, giving addition and scalar multiplication, completeness, and orthogonal projection tools (Axler, 2020).

The reason that Hilbert spaces are the ideal setting for this sort of model is due to this structure that they give to the underlying set and metric of the space. Within a Hilbert space, you have enough rigidity to claim things like fixed point sequences converge, which is incredibly important when using Landweber regularization based on Banach fixed point theorem (Axler, 2020; Engl et al., 1996). As such, there is a general consensus that Hilbert spaces are the correct setting for any economic or econometric work done, as seen in Carrasco et al. (2007) and Engl et al. (1996). Hence, developing an understanding of this underlying setting is important as a groundwork for understanding any model developed within this space.

7 Linear Inverse Problems

We now summarize the 2007 work of Marine Carrasco, Jean-Pierre Florens, and Eric Renault, who we will refer to as “the authors”, from their chapter in the *Handbook of Econometrics*, *Linear Inverse Problems in Structural Econometrics Estimation Based on Spectral Decomposition and Regularization* (Carrasco et al., 2007), which we will refer to as “the chapter”. We will do this through a brief introduction on the main points of the paper, and the exploration of four relevant examples to properly communicate the relevant details of the work.

We begin with our introduction of the work. The chapter is an exposition on the argument that we can validly proposition a variety of structural econometric questions as inverse problems to aid in the estimation of relevant economic information. This estimation, in some cases, is nontrivial: we see ill-posedness arise in many cases where either 1) the data or model is not well-specified, or 2) the problem is severely ill-conditioned in the sense of Hadamard, so small perturbations in the data can induce large changes in the solution (Carrasco et al., 2007). To address this, the authors also cover the topic of regularization and spectral decomposition to remove the noise of outliers and stabilize inversion, allowing for a much cleaner convergence within their estimation algorithms (Carrasco et al., 2007).

The way in which the authors argue this point comes in two forms, conveniently referenced as integral equations of the first and second kinds, respectively (Carrasco et al., 2007). The focus of the chapter is on recovering unknown functions that enter the model through expectations or integral operators (for example, conditional expectations or convolution kernels), hence the naming convention of "integral equations" (Carrasco et al., 2007).

Thus the argument that structural econometric questions can be posed as inverse problems comes in two operational forms, that of integral equations of the first and second kinds (Carrasco et al., 2007). Integral equations of the first kind come in the form $K\varphi = r$ (Carrasco et al., 2007). r represents a function (or vector/matrix of functions) constructed from real world data that we are interested in, such as conditional expectations or empirical moments; K represents some operator that is usually compact (Carrasco et al., 2007). Note that operators imply a linear map between Hilbert spaces in this context, and a compact operator maps a bounded set to a relatively compact set (Carrasco et al., 2007). φ, then, represents the unknown structural object we want to recover: an infinite-dimensional parameter such as a regression function, a structural response function, a density, or more generally a latent economic function living in an L^2 space (Carrasco et al., 2007).

Having explained the two different types of integral equations as the versions of inverse problems, we now look at both simple and more complicated examples for each type of integral equation, to exemplify the point of framing structural econometric problems in each way.

Beginning with integral equations of the first kind, we start with a simple example: a direct regression model (Carrasco et al., 2007). Define a model

$$Y = \varphi(Z) + U, \quad E[U \mid Z] = 0, \quad K\varphi = r. \tag{7.1}$$

To make this a concrete example, we want this model to symbolize a real-world economic relationship. So, define: Y: household monthly electricity consumption (in kWh); Z: household electricity price per kWh (a scalar); $\varphi(Z)$: the true average demand function mapping price Z to expected consumption; U: idiosyncratic shock to consumption (weather, household events) with mean zero given Z; K: the identity operator on functions of Z, so $(K\varphi)(z) = \varphi(z)$; r: the conditional mean of Y given Z, so $r(z) = E[Y \mid Z = z]$. By construction, the equation $K\varphi = r$ means that the true demand function φ equals the regression function $E[Y \mid Z]$ (Carrasco et al., 2007). We see that this is the "simple" example due to how K is defined: simply using the identity operator mapping φ to φ, we get the identity above ($\varphi(Z) = E[Y \mid Z]$) which allows us to just say that estimating φ is just nonparametric regression of electricity use on price (i.e. demand for electricity).

A more involved version of this model can be defined with the same basic setup (Carrasco et al., 2007). Define $Y = \varphi(Z) + U$, $E[U \mid W] = 0$, $K\varphi = r$, where: Y: individual labor supply (hours worked per week); Z: offered wage (endogenous, correlated with unobserved preferences); W: policy-driven variation, e.g. a randomized training assignment or eligibility instrument; $\varphi(Z)$: structural labor-supply function mapping wage Z to expected hours worked; U: unobserved preference and constraint shocks, mean zero given W; K: conditional expectation operator $K\varphi(w) = E[\varphi(Z) \mid W = w]$ (Carrasco et al., 2007); r: data-side conditional mean $r(w) = E[Y \mid W = w]$. This new construction gives us a way to solve a very different problem: $K\varphi = r$ means that taking the true structural labor-supply function and averaging it over the wage distribution using the conditional operator K, you will get the observed data r as a function of the wage-distribution policy W (Carrasco et al., 2007).

We now move to integral equations of the second kind, beginning with a simple smoothing example (Carrasco et al., 2007). Consider the model $(K\varphi)(x) = \int k(x,s)\,\varphi(s)\,ds$, $\varphi(x) - (K\varphi)(x) = r(x)$, which we can write compactly as $(I - K)\varphi = r$ (Carrasco et al., 2007). By construction, the equation $(I - K)\varphi = r$ means that at each date x, the difference between the desired current policy $\varphi(x)$ and the lingering effect of past policies $(K\varphi)(x)$ must exactly offset the exogenous pressure $r(x)$ (Carrasco et al., 2007). Solving the second-kind integral equation $(I - K)\varphi = r$ thus gives the policy rule φ that best balances current choices against lingering effects (Carrasco et al., 2007). Also note that $Y(x)$ is not needed explicitly here; φ itself is the object of interest rather than a constructed outcome variable.

We now turn to a more involved second-kind example with an additive structure (Carrasco et al., 2007). Consider the model $Y = \varphi(Z) + \psi(W) + U$, $E[U \mid Z, W] = 0$, which after algebra and conditioning can be rewritten as $(I - K)\varphi = r$ (Carrasco et al., 2007), where $K\varphi(z) = E\Big(E[\varphi(Z) \mid W] \Big| Z = z\Big)$ and $r(z) = E[Y \mid Z = z] - E\Big(E[Y \mid W] \Big| Z = z\Big)$ (Carrasco et al., 2007). Here K maps the policy response function φ into a function of Z that reflects how Z and W co-move through their joint distribution, and r is constructed entirely from observables to cancel the demographic component ψ (Carrasco et al., 2007). Solving this second-kind equation for φ gives an estimate of the structural policy effect of Z on healthcare spending, accounting for demographic biases $\psi(W)$.

Hence we see integral equations of the first kind being used for more standard estimation problems where the unknown appears only through K, whereas equations of the second kind are used for more convoluted settings where the unknown appears both directly and through K, often in the presence of nuisance components such as demographic bias and feedback effects (Carrasco et al., 2007). Ill-posed problems arise in practice once K is a genuinely compact operator built from conditional expectations or integral kernels; these are the cases for which the chapter's regularization machinery is developed (Carrasco et al., 2007). This is not how we will structure our models; instead we will define φ to be the economic state or information of interest (for example, a latent output gap), and K to be a compact forward operator based on structural econometric models that map that latent state into observables (Carrasco et al., 2007). '

8 Methodology

We now explain the underlying algorithm, objects relevant to said algorithm, and certain design choices in a general manner. Then we address a variety of challenges within our methodology, and certain cases or assumptions that must be noted. Lastly for clarity, we outline a specific example of estimating the output gap in the economy.

8.1 Model Outline

We begin by outlining our algorithm as follows. Initially, define $\varphi \in \mathcal{H}$ as the unknown economic state or parameter we wish to recover. Define K_1 as the forward operator implied by relevant structural econometric models. Lastly, define r as an object built from data or observables (macro series or transformations of macro series). This data should be the relevant output for the model(s) making up K_1. By construction, we have the following equation:

$$K_1 \varphi = r. \tag{8.1}$$

This is clearly the proper way to begin an inverse problem (Carrasco et al., 2007). This also allows us to say multiple things about the relationship between $K_1\varphi$ and r, which we will use later. Specifically note that this relationship is implied; the equation literally means that what the structural model says should be the state of the economy based on the true state of some economic variable is equal to the observed state of the economy, which holds trivially based on the descriptions of the given objects.

As the linear forward map from the economy's latent state φ to the model-implied observables r, K_1 is almost guaranteed to be ill-conditioned. This is for two main reasons. First, structural econometric models are too low-dimensional to take into account all relevant changes in the economy based on a shift of some economic variable, meaning that some small perturbation of φ might cause a large change in r due to some unobserved effect of that perturbation. Second, certain directions in the state space have only a weak effect on the data (Carrasco et al., 2007). In total, we cannot say that K_1 truly explains the comovement of φ and r, which is why we must treat this as an ill-posed problem and use regularization to estimate the true φ (Carrasco et al., 2007).

Now, given ill-posedness we must estimate φ by obtaining $\hat{\varphi}$ through some regularization schema. In our case we choose to do Landweber iteration, a type of gradient descent, because it allows us to use fixed-point-theorem language and logic (Engl et al., 1996). Some other options could have been Tikhonov or Ridge regularization (Engl et al., 1996), but those might be explored some other time.

Define an iteration as

$$\hat{\varphi}_{n+1} = \hat{\varphi}_n + \tau K_1^* \Big(r - K_1 \hat{\varphi}_n\Big), \tag{8.2}$$

with τ chosen such that the iteration is a contraction (Carrasco et al., 2007; Engl et al., 1996). We may set a step size τ as fixed a priori, or we can have a step size that is optimized in some way throughout the algorithm; either way, so long as the function is a contraction, the

Banach fixed point theorem states that we will converge (Axler, 2020).

One problem that arises here is the iteration at which we should stop iterating. We address this specific question in-depth later, but using the Morozov Discrepancy Principle (Engl et al., 1996; Morozov, 1984) we may find an iteration N^* based on the level of noise δ in the data such that $\hat{\varphi}_{N^*}$ is the best guess of φ based on data r (Anzengruber & Ramlau, 2010). Specifically, we calculate

$$N^* = \min \{n \in \mathbb{N} : \|r - K_1 \hat{\varphi}_n\| \leq \delta\} \tag{8.3}$$

to get our stopping parameter.

To perform our initial Landweber iteration, we need to choose an initial guess for $\hat{\varphi}_0$, which we can initialize as $\hat{\varphi}_0 = 0$ and still say that the limit of $\hat{\varphi}_n$ is the minimum-norm least-squares estimate of $K_1 \varphi = r$ by standard accounts (Engl et al., 1996). Initializing at 0 also gives us a clean guess, unbiased by any prior information. If we wanted to make this more accurate based on relevant statistics, we could also create a prior using the HP filter (Hodrick & Prescott, 1997) or the CBO potential output estimate (Shackleton, 2018); the theoretical guarantees of the algorithm are unaffected by this substitution, as the MDP stopping index N^* adjusts accordingly. Post this initialization, we perform iterations as above for $n = 0, \ldots, N^*$ to obtain a final guess $\hat{\varphi}_{N^*}$ for our first set of iterations.

Next, we define $\varepsilon = r - K_1 \hat{\varphi}_{N^*}$ as the residual of our estimation, with $\|\varepsilon\| \leq \delta$ by the Morozov Discrepancy Principle (Engl et al., 1996; Morozov, 1984). It represents all of the observables that were not captured by our structural model and estimation of φ, and hence can help us see where we fall short. We then want some way of incorporating this residual into our full model, as our model might not capture it for reasons like low-dimensionality or a lack of scope for explanation of data variability.

We address this in the following manner: define a new data matrix $Z \in \mathcal{H}_Z$, where $\mathcal{H}_Z$

is some Hilbert space. Z specifically is nonstandard economic data, because the residual we want to try to explain is explicitly not captured by the prior variables we took into account in our structural models. Using a kernel-based machine learning algorithm like a Neural Tangent Kernel (Jacot et al., 2018), we estimate the equation

$$\varepsilon = \psi(Z, \hat{\varphi}_{N^*}) + \text{err}, \quad E[\text{err} \mid Z] = 0, \quad \text{Cov}(Z, \text{err}) = 0 \tag{8.4}$$

with ψ being our kernel-based model. The reason we want to use machine learning here is because we are focusing on strictly accuracy; given that we want to debias as much as possible, we need a tool that will explain as much of the variability in the model residuals to debias our results (Jacot et al., 2018). We specifically train ψ on both Z and $\hat{\varphi}_{N^*}$ to allow our eventual operator to be well-constructed on φ.

Then, we need to construct our machine learning operator K_2. The way to do this varies based on the machine learning method chosen, with the most generic way being taking the Jacobian of ψ to get

$$K_2 \approx \left. \frac{\partial \psi}{\partial \varphi} \right|_{\hat{\varphi}_{N^*}}. \tag{8.5}$$

Using a Neural Tangent Kernel, we can simply extract the kernel from the model and use that as the compact linear operator (Jacot et al., 2018).

Lastly, we now define our complete operator $K = K_1 + K_2$. This construction makes sense, because we are adding in all of the effects that were not captured by the structural operator K_1 alone that were left in r, clearly shown through the way residuals are constructed: $\varepsilon = r - K_1\hat{\varphi}_{N^*}$, where $\|\varepsilon\| \leq \delta$ by the Morozov Discrepancy Principle (Engl et al., 1996; Morozov, 1966, 1984).

Then using K, we complete one more iteration of Landweber to estimate the true φ

based on our debiased operator. Initialize $\hat{\varphi}_0^{\text{debiased}} = \hat{\varphi}_{N^*}$, and iterate

$$\hat{\varphi}_{n+1}^{\text{debiased}} = \hat{\varphi}_n^{\text{debiased}} + \tau K^* \Big(r - K \hat{\varphi}_n^{\text{debiased}} \Big), \tag{8.6}$$

to find a new debiased estimation $\hat{\varphi}^{\text{debiased}} = \hat{\varphi}_{N_2^*}^{\text{debiased}}$, where N_2^* is the stopping place for the new debiased operator, calculated as

$$N_2^* = \min \left\{ n \in \mathbb{N} : \left\| r - K \hat{\varphi}_n^{\text{debiased}} \right\| \leq \delta \right\}. \tag{8.7}$$

Note the use of the operator K rather than K_1 within both the Landweber regularization and computation of the stopping parameter. In the trivial case where there is no residual and $K_1 \hat{\varphi}_{N^*} = r$, $N_2^* = 0$ as we do not need to take any new steps to get the optimal choice. In the complex case where $\delta > 0$ or K_1 misspecified, $K_1 \hat{\varphi}_{N^*} = r$ so $N_2^* > 0$. In general, $N_2^* \geq 0$.

Thus we estimate φ, the true latent state economic variable, through Landweber iteration and machine learning estimation, and obtain not only our guess $\hat{\varphi}_{N_2^*}^{\text{debiased}}$ but also information on the way this variable affects the economy in nonstandard ways in the form of K_2.

8.2 Addressals

It might be important to distinguish the finite and infinite-dimensional cases. The finite dimensional case is trivial: $\varphi \in \mathbb{R}^T$ is a finite vector and K_1, K_2, Z, and r are finite-dimensional matrices living in some Hilbert spaces (Axler, 2020). This case represents the empirical version of the problem, as it might be done in practice. One key note for the finite dimensional case lies in the Landweber regularization equation; in the general case we use

$$\hat{\varphi}_{n+1} = \hat{\varphi}_n + \tau K_1^* \Big(r - K_1 \hat{\varphi}_n \Big), \tag{8.8}$$

specifically with the adjoints K_1^* and K^* for each step. In the finite case, however, the adjoint is simply the transpose of K_1 and K. Thus we can say that for the finite case,

$$\hat{\varphi}_{n+1} = \hat{\varphi}_n + \tau K_1^\top \Big(r - K_1 \hat{\varphi}_n\Big), \tag{8.9}$$

and get the same effect (with the same for the second set of Landweber iterations using K).

The infinite-dimensional case is less trivial but still simple, representing the true functional-analysis version of the problem where φ is a single variable function and K_1, K_2 are bounded operators (often compact) acting on that function space (Carrasco et al., 2007). This represents the case where the Hilbert-space and inverse problem framework become important (Axler, 2020).

One other note is on how this model fits into the wider discourse of hybrid models discussed in the earlier portion of this paper. As explained, we want a model that estimates the true economic variable not just using structural features and machine learning models separately, but in conjunction in some way. This model satisfies this condition clearly; given the method of construction for $\hat{\varphi}^{debiased}$ in second Landweber step, we estimate the variable using both the structural and machine learning operators at the same time. This allows for the sequence $(\hat{\varphi}_n^{debiased})_{n=0}^{N_2^*}$ to balance itself between the structural portion and estimated machine learning portion in the same step, rather than optimizing for the structural part and then machine learning part iteratively. Moreover, given the machine learning kernel is trained on the residuals rather than the data in its entirety, the structural kernel will be the focal explanation point for our results, which is exactly what we want. Hence, this model addresses all major flaws within modern econometric models while retaining mathematical robustness within method and setting.

8.3 Applied Example

We now explain a relevant example to illustrate how our methodology might work in practice. In this example, the goal is to find the true output gap in the US economy, with output gap being defined as "the difference between an economy's actual output and its potential output" (International Monetary Fund, 2013). We measure this using the New Keynesian IS Curve (Galí, 2015), given by

$$x_t = E_t[x_{t+1}] - \frac{1}{\sigma}\Big(i_t - E_t[\pi_{t+1}] - r_t^n\Big), \tag{8.10}$$

where x_t is the output gap, defined as the difference between actual output and potential output (International Monetary Fund, 2013); $E_t[x_{t+1}]$ is the expectation at time t of next period's output gap; i_t is the nominal interest rate; $E_t[\pi_{t+1}]$ is expected inflation at time $t+1$; r_t^n is the natural real interest rate, i.e., the real rate consistent with zero output gap (Galí, 2015); and $\sigma > 0$ is the intertemporal elasticity of substitution.

To begin, stack a set of periods $t = 1, \dots, T$. Define φ as the true output gap of the US over time, such that $\varphi = (x_1, \dots, x_T)^\top$. Define r to be an object of data for the variables relevant to the New Keynesian IS Curve outlined above, at times $t = 1, \dots, T$. Lastly, define K_1 as the structural operator encoding the New Keynesian IS relationship, mapping the latent output gap φ to the variables in r at the quantities outlined by intertemporal output optimization (Galí, 2015). Outlining the actual structure of r and K_1, we can base these strongly on the NK IS curve:

$$r_t = \frac{1}{\sigma}\Big(i_t - E_t[\pi_{t+1}] - r_t^n\Big) \in \mathbb{R}, \tag{8.11}$$

$$r = \begin{pmatrix} r_1 \\ \vdots \\ r_T \end{pmatrix} \in \mathbb{R}^T, \tag{8.12}$$

$$K_1 = \frac{1}{\sigma}\begin{pmatrix} 1 & -1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & -1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & -1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 1 & -1 \\ 0 & 0 & \cdots & 0 & 0 & 1 \end{pmatrix} \in \mathbb{R}^{T\times T}. \tag{8.13}$$

We factor the entire K_1 matrix by $\frac{1}{\sigma}$ to scale by the inverse of the intertemporal elasticity of substitution (Galí, 2015). Thus for any given $t \in \{1, \cdots, T\}$, we have $(K_1\varphi)_t = \frac{1}{\sigma}(x_t - E_t[x_{t+1}]) = r_t$, validating the setup of our inverse problem.

Now beginning Landweber iterations, we initialize $\hat{\varphi}_0 = 0$ and iterate until we reach $\hat{\varphi}_{N^*}$. Using this estimator, we calculate the residuals $\varepsilon = r - K_1\hat{\varphi}_{N^*}$. We also define Z to be nonstandard financial and macroeconomic features that are not directly captured by the structural NK IS operator, such as BBB–AAA loan spreads, mortgage rates, VIX, and other financial stress indicators that may explain systematic variation in ε beyond what K_1 can account for (Carrasco et al., 2007).

We then train a neural tangent kernel on the residuals ε using Z and $\hat{\varphi}_{N^*}$, and then take the Jacobian of the network with respect to φ, evaluated at $\hat{\varphi}_{N^*}$, to construct K_2 (Jacot et al., 2018). Then, construct $K = K_1 + K_2$, initialize $\hat{\varphi}_0^{\text{debiased}} = \hat{\varphi}_{N^*}$, and do another round of Landweber, stopping at the Morozov Discrepancy Principle parameter

$$N_2^* = \min\left\{n \in \mathbb{N} : \left\| r - K\hat{\varphi}_n^{\text{debiased}} \right\| \leq \delta\right\} \tag{8.14}$$

to obtain $\hat{\varphi}_{N_2^*}^{\text{debiased}}$, the debiased estimator of the true latent output gap in the economy.

9 Morozov Discrepancy Principle

One problem arises with our algorithm, as mentioned prior: how do we validate the need for a machine learning step in the form of K_2? To claim that the reason we may not converge to a $\hat{\varphi}$ such that $r - K_1\hat{\varphi} \neq 0$ is due to a misspecification of the structural econometric models is a topic we discussed at length previously in this paper, but not in a technical manner. Hence, we must prove mathematically that there is a suitable reason to believe that the gap between $K_1\hat{\varphi}$ and r is not due to randomness in r, or that not enough steps of Landweber iteration have been done ($\hat{\varphi}$ is still misspecified).

The answer comes in the form of the Morozov Discrepancy Principle (Morozov, 1984), giving us a stopping parameter for Landweber iteration guaranteeing a best guess of $\hat{\varphi}$, even if our data r is noisy. The principle is explained in a relevant context in (Engl et al., 1996), but we explain an outline in simple terms here. Define a noise parameter δ as the degree to which our data r is noisy, i.e. a quantifier of how much the observed data in r differs from the true economy that we wish to measure. Suppose we perform Landweber iteration and stop at some iteration n. If it is the case[1] that $\|r - K_1\hat{\varphi}_n\| \leq \delta$, we know we have overfit our estimate.[2] The Morozov Discrepancy Principle (MDP) takes this realization to say that when doing Landweber iteration, the stopping parameter (i.e. the iteration we stop at) should be based off of the noise parameter δ (Morozov, 1984). Engl et al. (1996) defines their stopping parameter differently, but we will define our stopping parameter as the following:

$$N^* = \min\{n \in \mathbb{N} : \|r - K_1\hat{\varphi}_n\| \leq \delta\}. \tag{9.1}$$

We can describe N^* as the first iteration such that $\|r - K_1\hat{\varphi}_{N^*}\| \leq \delta$. By the MDP, we can conclude that $\hat{\varphi}_{N^*}$ is approximately the best possible estimation of the true economic variable

[1] As $n \to \infty$, $\|r - K_1\hat{\varphi}_n\| \to 0 \leq \delta$ by construction; hence there exist many $n \in \mathbb{N}$ such that $\|r - K_1\hat{\varphi}_n\| \leq \delta$, so the inequality is well-specified and there are cases in which this happens.

[2] We know we are overfitting because δ is defined to represent the noise floor, so fitting below it means we are fitting on measurement error or randomness.

we wish to recover (Engl et al., 1996)(Morozov, 1984). Given that we take the possible misspecification of r into account and using the MDP we converge to a good estimator $\hat{\varphi}$, we can verifiably say that any large residuals are the cause of a misspecification of K_1, the structural operator.

Remark 1. *δ is commonly specified either completely exogenously or as a manipulation of exogenous data. For this context, one can do something akin to (Chang & Li, 2015), where we compare results between methods of estimation or changes vintage-to-vintage for a given dataset to estimate the error in measurement. This might be done in the form of the following formula:*

$$\delta = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(r_t^f - r_t^a)^2}, \tag{9.2}$$

where r^f denotes the final version of the data and r^a denotes some prior version (Chang & Li, 2015). One might also be able to follow the steps outlined in (Jahn & Quinto, 2023) if one does not want to define a specific formula as above.

10 Conclusion

Solving the macroeconomic forecasting dilemma does not require choosing between tradition and innovation. Instead, it involves finding methods that bring together the best features of both approaches. The analysis in this paper shows that neither structural econometric models nor pure machine learning models alone can fully meet the demands of policy-relevant forecasting. Structural models excel in interpretability and credibility but falter in flexibility and crisis performance. ML models excel in capturing complex patterns but struggle with explainability and institutional trust.

Hybrid methodologies offer a promising compromise. When built with clear definitions and real functional interdependence between theory-driven and data-driven components, they can be modular in structure, adaptive in practice, and explainable enough for communication with policymakers (Batool, 2022; Buckmann et al., 2021; Kristjanpoller, 2024; Lee, 2025).

As an example, we developed a unique model derived from Carrasco et al. (2007) based on inverse problem theory, structural econometric operators,Landweber iteration, and machine learning models like neural tangent kernels to recover a debiased, accurate estimate of some latent economic state φ. In the future, this work might be developed by furthering the logic to apply to a broader context or expanding the scope of economic problems addressed. In addition, empirical testing and comparison is needed to concretely say how well this model performs in practice.

11 Acknowledgments

I sincerely thank Professor Bobby Wilson at the University of Washington Mathematics Department for his invaluable guidance and patience while advising me through this work.

References

Anzengruber, S. W., & Ramlau, R. (2010). Morozov's discrepancy principle for Tikhonov-type functionals with nonlinear operators. *Inverse Problems*, *26*(2), 025001. https://doi.org/10.1088/0266-5611/26/2/025001

Axler, S. (2020). *Measure, integration & real analysis* (Vol. 282). Springer. https://measure.axler.net

Batool, K. (2022). A hybrid model of machine learning and econometrics for volatility prediction: Evidence from the KSE-100 index. *Research Management Studies Journal*, *15*(3), 45–57. https://www.rmsjournal.com/index.php/admin/article/download/125/84

Buckmann, M., Joseph, A., & Robertson, H. (2021). *An interpretable machine learning workflow with an application to economic forecasting* (Working Paper Series). Bank of

England. https://www.bankofengland.co.uk/working-paper/2022/an-interpretable-machine-learning-workflow-with-an-application-to-economic-forecasting

Carrasco, M., Florens, J.-P., & Renault, E. (2007). Linear inverse problems in structural econometrics: Estimation based on spectral decomposition and regularization. In J. J. Heckman & E. E. Leamer (Eds.), *Handbook of econometrics* (pp. 5633–5751, Vol. 6B). Elsevier. https://doi.org/10.1016/S1573-4412(07)06077-1

Chang, A. C., & Li, P. (2015). Measurement error in macroeconomic data and economics research: Data revisions, gross domestic product, and gross domestic income. *Finance and Economics Discussion Series*, (2015-102). https://www.federalreserve.gov/econresdata/feds/2015/files/2015102pap.pdf

Engl, H. W., Hanke, M., & Neubauer, A. (1996). *Regularization of inverse problems*. Kluwer Academic Publishers.

Galí, J. (2015). *Monetary policy, inflation, and the business cycle* (2nd ed.). Princeton University Press.

Gietner, B. (2025). In defense of machine learning in economics. https://diddigest.substack.com/p/in-defense-of-machine-learning-in

Hammann, D., & Wouters, M. (2025). Explainability versus accuracy of machine learning models: The role of task uncertainty and need for interaction with the machine learning model. *European Accounting Review*, *35*(2), 343–375. https://doi.org/10.1080/09638180.2025.2463961

Hanke, M., Neubauer, A., & Scherzer, O. (1995). A convergence analysis of the Landweber iteration for nonlinear ill-posed problems. *Numerische Mathematik*, *72*(1), 21–37.

Hodrick, R. J., & Prescott, E. C. (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, *29*(1), 1–16.

International Monetary Fund. (2013). What is the output gap? *Finance & Development*, *50*(3). https://www.imf.org/external/pubs/ft/fandd/2013/09/basics.htm

Iskhakov, F. (2020). Machine learning and structural econometrics: Contrasts and synergies. https://fedor.iskh.me/assets/pdf/irs2020mlreview.pdf

Jacot, A., Gabriel, F., & Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks [arXiv:1806.07572 [cs.LG]]. *Advances in Neural Information Processing Systems*, *31*.

Jahn, T., & Quinto, E. T. (2023). Noise level free regularization of general linear inverse problems. *SIAM Journal on Numerical Analysis*, *61*(3). https://doi.org/10.1137/22M1506675

Kilic, R. (2025). *Linear and nonlinear econometric models against machine learning: Realized volatility prediction* (Finance and Economics Discussion Series No. 2025-061). Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2025.061

Kristjanpoller, W. (2024). A hybrid econometrics and machine learning based modeling of realized volatility of natural gas. *Financial Innovation*, *10*(1), Article 45. https://doi.org/10.1186/s40854-023-00577-0

Landweber, L. (1951). An iteration formula for Fredholm integral equations of the first kind. *American Journal of Mathematics*, *73*(3), 615–624.

Lee, C. (2025). Integrating econometrics and artificial intelligence to enhance economic forecasting. *Journal of Applied Economics & Forecasting*, *74*(3), 245–263.

Morozov, V. A. (1966). On the solution of functional equations by the method of regularization. *Soviet Mathematics Doklady*, *7*, 414–417.

Morozov, V. A. (1984). *Methods for solving incorrectly posed problems*. Springer-Verlag.

Pricepedia. (2025). From econometrics to machine learning: The challenge of forecasting. https://www.pricepedia.it/en/magazine/article/2025/10/06/from-econometrics-to-machine-learning-the-challenge-of-forecasting/

Shackleton, R. (2018). *Estimating and projecting potential output using CBO's forecasting growth model* (Working Paper No. 2018-03). Congressional Budget Office. https://www.cbo.gov/publication/53558

Smets, F., & Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, *97*(3), 586–606.

Woloszyn, J., & Bukowski, S. (2025). The impact of AI on economic modelling. *Economic Research Studies Journal*, *28*(1), 640–660.