

Discussion Paper

Deutsche Bundesbank
No 06/2023

Asset allocation with recursive parameter updating and macroeconomic regime identifiers

Milad Goodarzi

(Goethe University Frankfurt)

Christoph Meinerding

(Deutsche Bundesbank)

Editorial Board:

Daniel Foos
Stephan Jank
Thomas Kick
Martin Kliem
Malte Knüppel
Christoph Memmel
Hannah Paule-Paludkiewicz

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet <http://www.bundesbank.de>

Reproduction permitted only if source is stated.

ISBN 978-3-95729-937-6

ISSN 2749-2958

Non-technical summary

Research question

Strategies to allocate an investor's wealth to different asset classes (asset allocation) rest on models that describe the time series of the returns of these investments. Such models contain unobservable parameters and variables which need to be estimated from the data. The amount of data increases over time, suggesting to reestimate the unobservable quantities on a regular basis. Consequently, these quantities change over time. The literature has documented that this type of uncertainty can be very costly: it can generate long-lasting, quantitatively significant risks because it can amplify the impact of macroeconomic shocks. In this discussion paper, we study what kind of data is suitable for the estimation.

Contribution

We analyze portfolio strategies in a dynamic asset allocation exercise in which the determinants of the underlying return dynamics are reestimated every quarter. We study the optimization problem of a finite-horizon investor who can invest in four different asset classes (stocks, corporate bonds, commodities, treasury bills). The returns of these asset classes depend on an unobservable variable, which can switch between different states, so-called regimes. The investor estimates the dynamics of these regimes with US data, using either macroeconomic or financial data.

Results

We document that the portfolio performance can be improved substantially if the estimation relies on macroeconomic data instead of financial return data when determining the dynamics of regimes. The strategy relying on financial market data becomes very conservative in particular in the years after the financial crisis of 2008. We find several reasons for this. First, the estimated regimes are generally much more dispersed and extreme in terms of expected returns and return volatilities when we use market data for the regime identification. As a result, in particular after the financial crisis of 2008, a large fraction of wealth is invested in riskless treasury bills. Second, when using market data, the economy at times enters a new, previously unobserved regime (similar to the widespread notion of a so-called "Peso problem"), leading to significant revisions in the portfolio weights. Third, in line with the recent macro-finance literature, we find that the effect of rare events like the 2008 financial crisis on the regime estimates is very long-lasting.

Nichttechnische Zusammenfassung

Fragestellung

Strategien zur Aufteilung des Vermögens auf verschiedene Anlageklassen (Asset Allocation) beruhen auf Modellen, die den zeitlichen Verlauf der Renditen dieser Anlagen beschreiben. Solche Modelle enthalten unbeobachtbare Parameter und Variablen, die aus vorliegenden Daten geschätzt werden. Im Zeitablauf nimmt die Menge der vorliegenden Daten zu, so dass es sinnvoll ist, die unbeobachtbaren Größen regelmäßig neu zu schätzen. Die geschätzten Größen verändern sich somit im Zeitablauf. In der Literatur wurde bereits gezeigt, dass diese Art von Unsicherheit hohe Kosten verursachen kann. Sie kann die Auswirkungen makroökonomischer Schocks verstärken und langfristige, quantitativ bedeutsame Risiken erzeugen. Wir untersuchen in diesem Diskussionspapier, welche Art von Daten zur Schätzung geeignet ist.

Beitrag

Wir analysieren Portfoliostrategien im Rahmen eines dynamischen Asset Allocation Problems, in dem die Bestimmungsgrößen der betrachteten Renditen jedes Quartal neu geschätzt werden. Wir untersuchen dabei das Optimierungsproblem eines Investors mit endlichem Investitionshorizont, der in vier Anlageklassen investieren kann: Aktien, Unternehmensanleihen, Rohstoffe und risikolose Geldmarktpapiere. Die Renditen dieser Anlageklassen hängen von einer unbeobachtbaren Variable ab, welche zwischen verschiedenen Zuständen, sogenannten Regimen, hin- und herwechseln kann. Der Investor schätzt den Verlauf der Regime mit Hilfe von US-Daten. Dazu verwendet der Investor entweder makroökonomische oder Finanzmarktdaten.

Ergebnisse

Wir zeigen, dass die Performance des Portfolios stark verbessert werden kann, wenn zur Bestimmung der Regime makroökonomische anstelle von Finanzmarktdaten herangezogen werden. Die Strategie, die auf Finanzmarktdaten basiert, ist insbesondere in den Jahren nach der Finanzkrise von 2008 sehr konservativ. Wir finden mehrere Gründe dafür. Erstens sind die geschätzten Regime im Allgemeinen sehr viel unterschiedlicher und extremer hinsichtlich erwarteter Renditen und Renditevolatilitäten, wenn zur Regimeschätzung Marktdaten verwendet werden. Im Ergebnis führt dies dazu, dass insbesondere nach der Finanzkrise von 2008 ein Großteil des Portfolios in risikolose Geldmarktpapiere investiert wird. Zweitens kann es bei der Verwendung von Finanzmarktdaten passieren, dass zu einem bestimmten Zeitpunkt ein völlig neues Regime identifiziert wird (ähnlich dem in der Literatur bekannten „Peso Problem“), was große Portfolioumschichtungen nach sich zieht. Drittens zeigen wir im Einklang mit der jüngsten makrofinanziellen Literatur, dass der Effekt seltener Ereignisse wie der Finanzkrise von 2008 auf die geschätzten Regime sehr langanhaltend ist.

ASSET ALLOCATION WITH RECURSIVE PARAMETER UPDATING AND MACROECONOMIC REGIME IDENTIFIERS

Milad Goodarzi*

Christoph Meinerding[‡]

December 12, 2022

Abstract: This article studies long-horizon dynamic asset allocation strategies with recursive parameter updating. The parameter estimates for the regime-switching dynamics vary as more and more datapoints are observed and the sample size increases. In such a setting, the globally optimal portfolio strategy cannot be determined due to computational complexity. Among a set of suboptimal strategies, the portfolio performance can be improved substantially if the dynamics of the regimes are estimated from fundamental macroeconomic instead of financial return data. Especially after highly uncertain times, the estimation based on financial market data identifies extreme regimes, leading to extreme hedging demands against regime changes.

Keywords: Regime switching models, asset allocation, macro-based portfolio strategies, parameter updating

JEL: G11, D83, E44

*Goethe University Frankfurt, Faculty of Economics and Business, E-mail: milad.goodarzi@hof.uni-frankfurt.de

[‡]Deutsche Bundesbank, Research Centre, E-Mail: christoph.meinerding@bundesbank.de

Our special thanks go to Fabio Girardi, Christian Schlag, Julian Thimme, Rüdiger Weber and colleagues and seminar participants at Goethe University and Deutsche Bundesbank for many helpful discussions over the course of this project. This work represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

When implementing a dynamic asset allocation strategy over a long horizon, an investor faces the key challenge that the estimates for the parameters of the data-generating process vary as more and more datapoints are observed and the sample size increases. In their seminal contribution, Collin-Dufresne, Johannes, and Lochstoer (2016) document that parameter learning can be very costly: it generates long-lasting, quantitatively significant macroeconomic risks because it amplifies the impact of macroeconomic shocks on agents' marginal utility substantially.

In this paper, we analyze portfolio strategies in a dynamic asset allocation exercise in which the parameters of the underlying return dynamics are updated recursively every period. In particular, we document that the performance of such strategies can be improved substantially if the estimation relies on fundamental macroeconomic data instead of financial return data when determining the dynamics of regimes. More precisely, we study the optimization problem of a finite-horizon CRRA investor who can invest in four different asset classes (stocks, corporate bonds, commodities, T-Bills) whose returns follow standard Markov switching VAR dynamics. The regimes are unobservable and the investor uses US post-war data – either macroeconomic or financial data – to estimate them. Parameter updating is incorporated in this setup by iterating the estimation and portfolio optimization on a quarterly basis with increasing sample size as more and more datapoints are observed. Our findings show that, especially after highly uncertain times like the burst of the dotcom bubble or the 2008 financial crisis, the estimation based on financial market data identifies extreme regimes and leads to a significant underperformance of the respective investment strategy.

To understand this key finding, it is important to recognize that, in a setup like ours, the *globally optimal* portfolio strategy cannot be determined. For one, we suffer from the curse of dimensionality as already highlighted by Barberis (2000). The estimated fundamental dynamics change with every newly observed datapoint, i.e. the parameters of the data-generating process themselves represent state variables whose evolution needs to be taken into account when determining the hedging demand of the investor. For another, the dynamics of these additional state variables are unknown, since the distribution from which new data observations are drawn every period is generally unknown. As a result, any numerical solution, e.g. through standard backward induction, requires simplifying (parametric or non-parametric) assumptions about these dynamics, which makes the numerical solution globally suboptimal to begin with.

Thus, in a setup with recursively updated parameters like ours, it is only possible to make

meaningful simplifying assumptions and then compare the resulting suboptimal solutions. In this vein, we make the following assumptions: we take into account the hedging demand of an investor with respect to regime changes (for a given set of parameter estimates). But, throughout the paper and as in Barberis (2000), we exogeneously set the hedging demand with respect to parameter changes to zero. It is intuitive to assume that an investor can grasp the former hedging demand more easily than the latter.

We then compare recursive portfolio strategies that differ by means of the data that is used for estimating one particular piece of the Markov switching VAR dynamics for returns, namely the underlying regimes. Our portfolio choice exercise generally follows Guidolin and Timmermann (2007), but with a few notable differences. In order to tease out the effect of different regime identifiers, we split the Markov switching VAR estimation procedure into two parts: (i) the regime identification using standard maximum likelihood estimation and (ii) the VAR estimation conditional on each regime that has been identified in Step (i). For the “macro-based” portfolio strategy, the regime identification (i) is executed with US consumption growth and inflation data. For the “market-based” strategy, we use stock and bond returns instead, as in Guidolin and Timmermann (2007). Once the regimes have been identified, Step (ii) is executed with asset return data for both strategies. Importantly, we deliberately do not use any predictor variables in Step (ii), so as to downplay any potential side effects from the predictability of financial returns by macroeconomic variables and focus only on the properties of the identified regimes.

Given the parameters of the Markov-switching VAR, we compute the optimal portfolio weights through standard backward induction, taking into account the hedging demand for regime changes and assuming Bayesian learning of the hidden regime by the investor, but disregarding the hedging demand for parameter changes as in Barberis (2000). In order to take the perspective of a long-horizon investor who wants to invest for 30 years from 1990 until 2020, we iterate the procedure forward quarterly, i.e. we augment the sample by the next data observation, repeat the estimation, reduce the investment horizon by 1 quarter, repeat the backward induction and so on and so forth. As time passes, more and more information becomes available and the investment strategy is readjusted every quarter, also taking into account the shrinking investment horizon as we approach the final year 2020.

As the key result of our paper, we find that the macro-based strategy outperforms the market-

based strategy substantially. It yields a 5.7% higher (annualized) certainty equivalent return, and the Sharpe ratios of the two strategies differ by 0.46. A first glance at the portfolio weights reveals that the market-based strategy becomes very conservative in particular in the years after the financial crisis of 2008, which can be traced back to different estimates of the regimes going forward. We find several reasons for this. First, the estimates reveal that, when using market data for the regime identification, the economy seems to enter a new regime that has never been observed before (similar to the widespread notion of a so-called “Peso problem”), leading to significant revisions in the portfolio weights. Second, we find that the estimated regimes are generally much more dispersed in terms of expected returns and return volatilities when we use market data for the regime identification. Moreover, the worst possible regime (by means of expected returns) is much more extreme and also more likely for the market-based estimation. We show that these features of the estimation amplify the hedging demand with respect to regime changes, so that, in particular after the disruptive financial crisis of 2008, the market-based portfolio strategy becomes very conservative, investing a large fraction of wealth in T-Bills. Third, we find that the effect of rare events like the 2008 financial crisis on the regime estimates is very long-lasting. This is in line with general idea of costly parameter learning brought up by Collin-Dufresne et al. (2016). More recently, Kozlowski, Veldkamp, and Venkateswaran (2020) have made a similar point, emphasizing the potential stickiness of belief changes after rare events.

We corroborate this key result with a set of robustness checks, for instance changes in the sample periods, different macroeconomic time series for the estimation, a combination of both macro and market data in the estimation, or different degrees of transaction costs. Overall, we conclude that, in a situation where the globally optimal solution to a portfolio problem is unknown because of parameter learning, macro data can enhance the risk-adjusted performance of the portfolio strategy.

The rest of the paper proceeds as follows. Section 1 reviews the related literature. In Section 2, we outline the estimation and the portfolio optimization procedure. Sections 3 and 4 present our main finding and a set of additional exercises to understand the mechanism behind it. Robustness checks are discussed in Section 5. Section 6 concludes.

1 Literature review

The issue of parameter learning is of key importance in studying risk premia in partial and general equilibrium models. In their seminal paper, Collin-Dufresne et al. (2016) study learning in an endowment economy and show that parameter uncertainty is an important source of long-run risk that can lead to very high equilibrium risk premia. Brennan and Xia (2001) analyze uncertainty about expected dividend growth in a general equilibrium model which can match historical stock price volatility and the equity premium. In the same vein, parameter learning has recently been used in partial equilibrium models as an amplification mechanism for the pricing of shocks (Lewellen and Shanken (2002), Weitzman (2007), Wang (2009), Peijnenburg (2018)). Our paper emphasizes the effect of uncertainty and recursive updating of parameters on long-horizon asset allocation and shows that we can robustify the performance by using macro data for the parameter estimation. Similarly, Brennan (1998) and Xia (2001) study the effect of uncertainty and learning about mean returns on the optimal investment strategy of an investor with long investment horizon, but in a much simpler setting, in which the optimal solution can be computed in closed form. Barberis (2000) studies optimal long-horizon portfolios with parameter uncertainty when returns are predictable.

Our paper also relates to a rich strand of literature featuring Hidden Markov Models for market returns or macroeconomic variables. Turner, Startz, and Nelson (1989) assume a Markov regime switching process for excess return volatility and document a negative correlation between variance and excess return. Relatedly, Brennan, Schwartz, and Lagnado (1997) and Della Corte, Sarno, and Tsiakas (2010) develop a trading strategy based on a Hidden Markov Model with penalizing switches between the states. Song, Wang, and Yang (2014) study the optimal dynamic asset allocation strategy when there is stochastic variation in stock returns. Markov switching models for inflation and GDP have been studied, among others, by Kim and Nelson (1999) and Fallahi (2011) who identify structural breaks in US GDP growth data. Barsoum and Stankiewicz (2013) and Nalewaik (2015) use a Markov regime switching model for forecasting GDP growth and inflation, respectively. Amisano and Fagan (2013) construct an early warning indicator for switches in the inflation regime. Klein and Shambaugh (2015) and Zhou (2020) use Markov regime switching models of inflation and growth for monetary policy analysis. Dergunov, Meinerding, and Schlag (2022) show that extreme inflation observations can help identify time variation in expected consumption growth and thereby help to understand

time variation in the correlation of stock and bond returns.

Our findings complement those of papers which study regime based portfolio optimization (Ang and Bekaert (2002), Ang and Bekaert (2004), Hess (2006), Guidolin and Timmermann (2007), Tu (2010) and Kole and Dijk (2017) among others). Typically, researchers use market returns for identifying the regimes; for example Honda (2003) documents the importance of considering regimes for mean returns, in particular for the hedging demand and very long investment horizons. Bulla, Mergner, Bulla, Sesboue, and Chesneau (2011) study the out-of-sample performance of a regime-based asset allocation and show that this strategy is profitable even after considering transaction costs. Pettenuzzo and Timmermann (2011) evaluate the effect of ignoring structural breaks in parameters of return predictability for long-term investment strategies. A group of papers (like Kritzman, Page, and Turkington (2012), Ilmanen, Maloney, and Ross (2014) and Kollar and Schmieder (2019)) also study the effects using macroeconomic variable for regime switching in asset allocation, but they disregard the aspect of parameter learning.

Our paper also relates to the literature on suboptimal, approximative solutions for complex optimization problems. The existing literature often uses grid-based dynamic programming, Markov chain approximations (MCAs) or finite difference methods for the Hamilton Jacobi Bellman (HJB) equation to find the globally optimal portfolio. But these methods are applicable only to low-dimensional problems (Brennan et al. (1997), Munk (2000), Munk and Sorensen (2010)). Some researchers have provided new methods to deal with the curse of dimensionality and find a suboptimal solution close to the global optimum in an efficient way (Bick, Kraft, and Munk (2013), Beard, Saridis, and Wen (1998), Chen and Jagannathan (2008) and Zeman (2010)). We face a dimensionality problem because of the parameter learning setup. We document that using a robust set of variables for parameter estimation helps in finding a better (though still suboptimal) portfolio allocation. In a sense, our paper is thus also related to the paper of Martin and Nagel (2022) about high-dimensional prediction problems in asset pricing.

Finally, our paper is also related to the enormous amount of literature which documents or neglects the existence of unspanned macroeconomic risks in bond yields and returns. Influential contributions have been made, among many others, by Bauer and Hamilton (2017, 2018), Coroneo, Giannone, and Modugno (2016), Joslin, Pribsch, and Singleton (2014), Chernov and Mueller (2012), Bikbov and Chernov (2010), Ludvigson and Ng (2009), Johannes, Polson, and Stroud (2009), and Ang and Piazzesi (2003). Building on this literature, de Pooter, Ravazzolo,

and van Dijk (2007) document the predictive power of macroeconomic variables for the term structure of interest rates in a setting with parameter uncertainty. We do not take a stand in this very broad debate, but we use inflation and real growth in our benchmark setup because these are the two most prominent variables in this literature.

2 Regime-based asset allocation strategies

The implementation of each investment strategy requires three steps: (i) identify the regimes, (ii) estimate the return dynamics of different asset classes conditional on each regime, and (iii) optimize the portfolio based on the estimated return dynamics. The portfolio optimization procedure (iii) follows mostly Guidolin and Timmermann (2007), but the estimation steps (i) and (ii) are markedly different. Guidolin and Timmermann (2007) execute these two steps jointly, using the same data for regime identification and return dynamics. We, however, want to tease out the effect of different regime identifiers for the performance of the portfolio. We therefore disentangle the two steps throughout the paper. We will discuss each step in the following subsections.

The three steps are carried out recursively, using expanding windows of quarterly data. The first iteration is done with data from 1948Q1 until 1989Q4, the second iteration is with data from 1948Q1 until 1990Q1 and so on, until 2020Q2.

We do so in order to take the perspective of a long-horizon investor who wants to invest for 30 years from 1990 until 2020. As time passes, more and more information becomes available and the investment strategy is recalibrated every quarter, also taking into account the shrinking investment horizon as we approach the final year 2020.

As will become clear below, the globally optimal solution to the portfolio choice problem under parameter learning is unknown, reflecting the curse of dimensionality. Therefore, all investment strategies discussed in this paper reflect approximations of this unknown optimal solution.

2.1 Identification of regimes

The first step is the identification of the regimes. The labeling of the two main investment strategies discussed in the following (“macro-based” or “market-based”) reflects the data that is used in this first step, namely macroeconomic time series (consumption growth and inflation)

or return data (stock and bond indices). In robustness checks, we will also analyze combinations of the two.

Let x_t be the vector of variables that are used for the regime identification. We assume a constant variance-covariance matrix and only allow for time-varying drifts. We estimate the following time series model

$$\Delta x_{t+1} = \mu(Z_t)\Delta t + \sigma\varepsilon_{t+1} \quad (1)$$

where the ε are i.i.d. standard normal shocks and Z_t is the current value of a k -state Markov chain. The estimation is done via the standard expectation-maximization (EM) algorithm based on Hamilton (1990). Following Guidolin and Timmermann (2005, 2007), we assume that the Markov chain has four states throughout the whole paper for the sake of comparability and tractability.

At each iteration, the Markov chain transition probability matrix P is stored, to be used in the simulation part later on. Moreover, we store the estimated regime switches. For each point in time, the algorithm delivers the ex-post probability that the economy was in a particular regime at a particular point in time, which is almost always very close to 0 or 1. We round these probabilities to 0 or 1 and thus assign one of the four regimes to each point in time. This leaves us with a time series of regimes for the economy. Based on these identified regimes, we divide our sample into four subsamples, i.e. one subsample for each regime. The parameters μ and σ from the first estimation step are not used in the following.

2.2 Estimation of return dynamics conditional on regimes

Having estimated a time series of historical regimes, we next estimate the dynamics of returns conditional on a given regime. Let y_t be the vector of quarterly returns of a set of assets (in the benchmark case a broad stock index, a long-term corporate bond, a short-term government bond, and a broad commodity futures index). We estimate the following VAR model on each subsample $j = 1, \dots, 4$ identified in the previous subsection

$$y_{t+1}^{(j)} = A^{(j)}y_t + \mu^{(j)}\Delta t + \Omega^{(j)}\varepsilon_t \quad (2)$$

where ε_t are i.i.d. standard normal shocks.¹ Combined with the transition probabilities from the previous subsection, we thus end up with Markov switching VAR dynamics for returns.

Our estimation strategy deviates from Guidolin and Timmermann (2007) in two important aspects. First, these authors combine the regime identification and the VAR estimation in one single estimation and maximize the joint likelihood function, whereas we separate the two steps in order to tease out the effects of different regime identifiers. The parameters and states that we estimate through our ad-hoc two-step procedure will not maximize the joint likelihood function globally. An idea to get closer to the global optimum would be a recursive procedure that loops over the two steps multiple times. We refrain from such a procedure because the scope of our paper is different. As outlined previously, we do not aim at finding the globally optimal portfolio strategy for our problem. Instead, we want to compare portfolio strategies relying on different regime identifiers with each other. For this purpose, we regard a lean and simplified estimation approach as more suitable.

Second, while Guidolin and Timmermann (2007) allow for additional predictor variables like the dividend yield in the VAR, we shut this channel down for clarity of exposition and run the VAR estimation conditional on the regime without any further predictor variables. We do so in order to avoid any potential side effects due to the predictability of financial returns by macroeconomic variables. The focus of our paper lies on the properties of regimes identified via macroeconomic or market data like, e.g., the differences in myopic and hedging demand that these different regime dynamics induce.

2.3 Portfolio choice based on estimated return dynamics

As the final step, we compute the optimal portfolio of a CRRA investor with utility defined over terminal wealth, W_{t+T} , coefficient of relative risk aversion γ , with T quarters investment horizon who is faced with the Markov-switching VAR dynamics for returns. In each iteration, we will assume that the final date is 2020Q3, i.e. the investment horizon T shrinks by one quarter in every iteration:

$$u(W_{t+T}) = \frac{W_{t+T}^{1-\gamma}}{1-\gamma} \quad (3)$$

¹In principle, forming subsamples by concatenating all datapoints that are assigned to one particular regime could be problematic for the estimation, in particular when the VAR allows for lags of y_t . However, we verify that our portfolio choice results are robust to this assumption by also running an exercise where we set $A = 0$. The results are discussed in Section 5.5.

The investor is assumed to maximize expected utility by choosing at time t a portfolio of stock, bond, commodity futures and riskless T-bills. Throughout the paper, we assume a relative risk aversion of $\gamma = 5$.

The investor knows her investment horizon and the Markov-switching nature of the returns, i.e. the optimal portfolio consists of a myopic demand and a hedging demand addressing the risk of future regime changes. However, in order to circumvent the curse of dimensionality outlined above, we assume that she disregards the parameter learning structure of the problem as in Barberis (2000). I.e., the optimal portfolio at time t is computed *as if* the parameters estimated at time t were constant. In this sense, the “optimal” strategy derived in this section is only an approximation to the unknown, globally optimal strategy which would also comprise a hedging demand for the risk of future parameter changes due to the iterative structure of the problem.

Consistent with common practice, e.g. Guidolin and Timmermann (2007), we rule out short selling. The investor’s optimization problem then is:

$$\begin{aligned} \max_{\omega_t} \quad & E_t \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right] \\ \text{s.t.} \quad & W_{t+1} = W_t (\omega_t' \exp(r_{t+1})) \\ & \omega_t \in [0, 1] \end{aligned} \tag{4}$$

ω_t is the vector of optimal portfolio weights at time t . r_t is the vector of asset returns.

Under power utility the value function simplifies to

$$J(W_t, r_t, \theta, \pi_t, T - t) = \frac{W_t^{1-\gamma}}{1-\gamma} Q(r_t, \theta, \pi_t, T - t) \tag{5}$$

where θ contains the parameters of the Markov switching VAR estimated in the previous step. π_t is the 4×1 vector of subjective probabilities for each of the four possible states conditional on information at time t . These probabilities serve as state variables.

Their dynamics follow from Bayesian learning

$$\pi_{t+1}(\hat{\theta}) = \frac{(\pi_t(\hat{\theta})\hat{P}_t)' \odot \eta(r_t; \hat{\theta})}{[(\pi_t(\hat{\theta})\hat{P}_t)' \odot \eta(r_t; \hat{\theta})]'\iota_k} \tag{6}$$

where $\odot$ denotes the element-by-element product, ι_k is a 4×1 vector of ones, and $\eta(r_t; \hat{\theta})$ is the

4×1 vector whose j -th element gives the density of the observable variable r_t in the j -th state at time t conditional on $\hat{\theta}$:

$$\eta(r_t; \hat{\theta}) \equiv \begin{bmatrix} (2\pi)^{-k/2} |\hat{\Omega}_1^{-1}|^{1/2} \exp \left[-\frac{1}{2}(r_t - \hat{\mu}_1 - \hat{A}_1 r_{t-1})' \hat{\Omega}_1^{-1} (r_t - \hat{\mu}_1 - \hat{A}_1 r_{t-1}) \right] \\ (2\pi)^{-k/2} |\hat{\Omega}_2^{-1}|^{1/2} \exp \left[-\frac{1}{2}(r_t - \hat{\mu}_2 - \hat{A}_2 r_{t-1})' \hat{\Omega}_2^{-1} (r_t - \hat{\mu}_2 - \hat{A}_2 r_{t-1}) \right] \\ \vdots \\ (2\pi)^{-k/2} |\hat{\Omega}_k^{-1}|^{1/2} \exp \left[-\frac{1}{2}(r_t - \hat{\mu}_k - \hat{A}_k r_{t-1})' \hat{\Omega}_k^{-1} (r_t - \hat{\mu}_k - \hat{A}_k r_{t-1}) \right] \end{bmatrix} \quad (7)$$

The investor does not know the current state of the Markov chain. She takes uncertainty about the current state into account in her portfolio decision. In order to determine the optimal portfolio, we simulate paths of the economy using the current Markov switching VAR parameters, and we assume that the investor updates her probabilities of the current state on these paths when new return observations come in. In other words, the investor assumes that she is going to update the probabilities of the states based on new return information, while sticking to the parameters of the Markov-switching VAR that have been estimated in the beginning.

Knowing that $Q(r_T, \theta, \pi_T, 0) = 1$ at time T , we can solve the optimization problem recursively backwards in time applying the Bellman principle:

$$Q(r_t, \theta, \pi_t, T-t) = \max_{\omega_t} E_t \left[\left(\frac{W_{t+1}}{W_t} \right)^{1-\gamma} Q(r_{t+1}, \theta, \pi_{t+1}, T-t-1) \right] \quad (8)$$

The solution is computed numerically in several steps: (i) discretize the state space, (ii) for every grid point of the state space, simulate returns from the estimated Markov-switching VAR dynamics, (iii) for every simulated return path, compute the updated path of probabilities π_{t+1} , (iv) compute the expectation in (8). As the data are quarterly, the optimization is also based on a quarterly discretization and the simulation step within the optimization is also quarterly.

Finally, when computing the portfolio performance, we also account for transaction costs caused by the quarterly rebalancing. Consistent with common practice in the asset allocation literature, we assume transaction costs proportional to the traded dollar amount, so that the portfolio performance remains independent of the size of the portfolio. For the baseline case we set transaction costs per trade equal to 0.5% of the traded dollar amount (equally for all assets), but we present results for different values in the robustness section below.

3 Main results

In this section we display the main results that we obtain for the baseline case. Further cases will be discussed in the subsequent sections, in order to illustrate the mechanism.

3.1 Data

All data used in this paper are quarterly US data, ranging from 1948Q1 to 2020Q3. For stock returns, we use continuously compounded S&P 500 index returns (month-end values) including dividends, obtained from CRSP. For the risk-free interest rate, we use the 1-month Treasury Bill rate. Long-term corporate bond returns are taken from the Ibbotson Stocks, Bonds, Bills and Inflation Yearbook. For commodities, we use the returns on the Thomson Reuters Equal Weight Continuous Commodity Index. The index is an equal weighted index of 17 commodities plus theoretical interest earned on the cash required to collateralize the futures trading. Inflation is the Consumer Price Index (All Urban Consumers) from the Bureau of Labor Statistics, consumption growth is obtained from the Personal Consumption Expenditures of the Bureau of Economic Analysis, both retrieved via FRED. Real GDP growth is taken from the OECD.

3.2 Results

For the baseline case, we use the longest possible data sample (from 1948Q1 to 2020Q3). The first iteration within the procedure outlined in the previous section is done with data from 1948Q1 to 1989Q4, i.e., the first portfolio weight that we compute is for 1990Q1. The investment horizon thus equals 123 quarters, and the investor can rebalance the portfolio every quarter. For computing the portfolio weight of the next period (i.e. 1990Q2), we use the data from 1948Q1 to 1990Q1, but the time horizon will be reduced by 1 quarter and so forth.

As regime identifiers, we use consumption growth and inflation for the macro-based strategy and stock and bond returns for the market-based strategy. Apart from the regime identifiers, everything else is the same for these two strategies. The investable asset classes are stocks, bonds, commodities and cash (risk-free investment). Table 1 reports the main results for the baseline case.

First of all, the Sharpe ratio is 0.66 for the macro-based strategy and 0.20 for the market-based strategy and the Sharpe ratio difference of 0.46 is statistically significant, according to

Table 1: **Baseline case performances**

	Macro-based strategy	Market-based strategy
Sharpe ratio	0.66	0.20
(p-value of the difference)	-	(0.04)
CEQ return	5.2%	-0.5%
Total return	7.3%	1.7%
VaR	15.5%	18.7%
Turnover	0.62	0.33
Skewness	-1.24	-1.58
Kurtosis	4.00	6.45

The table reports the performances of the different strategies in the baseline case. The p-value of the Sharpe ratio difference is computed using the Ledoit and Wolf (2008) methodology. CEQ is the annualized certainty equivalent return, Total return shows the annualized return of the strategy over the whole period. VaR refers to the 5% value at risk and turnover captures the changes in the portfolio weights across different asset classes and over time.

the statistical test for difference in Sharpe ratios proposed by Ledoit and Wolf (2008).² The investor in our model is not maximizing the Sharpe ratio, but expected utility, therefore it is theoretically more meaningful to compare certainty equivalents (CEQ):

$$CEQ = [e^{(1-\gamma)\mu + \frac{1}{2}(1-\gamma)^2\sigma^2}]^{\frac{1}{1-\gamma}} \quad (9)$$

For ease of exposition we report annualized certainty equivalent returns. As the table documents, the macro-based strategy also performs better in terms of certainty equivalents.

The table also shows that not only the Sharpe ratio is higher for the macro-based strategy, but also the total return, defined as the annualized return of the strategy over the whole investment period without adjusting for volatility. The outperformance is also visible in the value-at-risk. Here we report the 5% VaR, i.e. the (negative of the) 5% quantile of the quarterly portfolio return time series. This number can be interpreted as the maximum percentage loss over the next quarter which the investor expects with 95% confidence, when she follows the respective investment strategy. Similarly, the return of the macro-based strategy is also less negatively

²We thank Olivier Ledoit and Michael Wolf for providing the code on their website.




Figure 1: **Baseline results - cumulative returns**

This figure shows cumulative log returns of the different strategies in the baseline case with transaction cost.

skewed and less leptokurtic than the return of the market-based strategy.

Although our model explicitly accounts for transaction costs, we also report the portfolio turnover. Following DeMiguel, Garlappi, and Uppal (2009), we compute turnover as the average sum of the absolute values of all trades across the N assets:

$$\text{Turnover} = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N (|\hat{\omega}_{i,t+1} - \hat{\omega}_{i,t}|) \quad (10)$$

where $\hat{\omega}_{i,t+1}$ is the portfolio weight in asset i at time t . Importantly, the macro-based outperforms the market-based strategy although it has a higher turnover, indicating that the outperformance is not simply due to transaction costs. For the baseline case we set transaction costs equal to 0.5% per trade (equally for all assets), but we present results for different values in the robustness section below.

The main observations are confirmed by Figure 1 and 2, which show the portfolio performance graphically. For the first few years, both strategies move pretty much in parallel. Then, in the late 1990s, there is a parallel shift downwards in the market-based strategy. Both strategies see an almost similar loss around 2002 and around 2008. From 2008 onwards, both strategies diverge substantially. While the macro-based strategy continues to show a good performance, the market-based strategy becomes almost flat until 2018.




Figure 2: **Baseline results - quarterly returns**

This figure shows quarterly returns of the different strategies in the baseline case with transaction cost.

A look into the portfolio holdings reveals that the investor following the market-based strategy invests nearly 100% of his wealth into the risk-free asset after 2008, while the investor following the macro-based strategy continues investing in risky assets like stocks and commodities. This key results survives many robustness checks, as we show below. In order to understand it, we will analyze a few instructive subcases of the general optimization problem in the following section.

4 Understanding the mechanism

In the previous section we have shown that the macro-based strategy outperforms the market-based strategy along various measures. In this section we shed light on why this is the case.

4.1 Myopic vs. hedging demand

We study portfolio optimization over a very long time horizon (up to 30 years). For such long-horizon portfolios, the dynamics of the investment opportunity set are key. As Merton (1971) has pointed out, the optimal portfolio of a risk-averse investor who is confronted with a time-varying investment opportunity set satisfies a three-fund separation principle. It is composed




Figure 3: **Myopic case**

This figure shows the cumulative log returns of the different strategies in the myopic case with transaction cost.

of an investment in the risk-free asset, an investment in the tangency portfolio (the so-called myopic portfolio), and an investment in the so-called hedge portfolio.

The composition of this hedge portfolio depends on the covariance between asset returns and state variables determining the investment opportunity set. For example, if the return of a risky asset like a stock index is negatively correlated with changes in its volatility, a risk-averse investor will hold a smaller position in this risky asset than an otherwise identical myopic investor, reflecting the fact that negative returns are associated with increasing volatility, i.e., deteriorating investment opportunities going forward. There is no hedging demand if the investment opportunity set is time-invariant or if the investment opportunity set is uncorrelated with asset returns.

Since we solve for the optimal portfolio numerically, we cannot decompose the portfolio in closed form. Therefore we follow a slightly different strategy. We run the same computations as in the baseline case except for one major change. In the portfolio optimization step (Equation (4)), we set $T = t + 1$, i.e. we assume an investment horizon of 1 quarter in every iteration of the loop. The resulting portfolio strategy is purely myopic. Figure 3 and Table 2 show the performances of the two respective strategies.

As can be seen from Table 2, the outperformance of the macro-based strategy is reduced when the investor acts myopically. For instance, the difference in certainty equivalent returns shrinks

Table 2: **Myopic case performances**

	Macro-based strategy	Market-based strategy
Sharpe ratio	0.61	0.34
(p-value of the difference)	-	(0.03)
CEQ return	4.7%	0.9%
Total return	7.4%	3.6%
VaR	17.8%	19.7%
Turnover	0.26	0.63
Skewness	-0.69	-1.3
Kurtosis	1.34	2.4

The table reports the performance of the different strategies in the myopic case. The p-value of the Sharpe ratio difference is computed using the Ledoit and Wolf (2008) methodology. CEQ is the annualized certainty equivalent return, Total return shows the annualized return of the strategy over the whole period. VaR refers to the 5% value at risk and turnover captures the changes in the portfolio weights across different asset classes and over time.

from 5.7% to 3.8% annually. The table suggests that this shrinkage is exclusively due to a better performance of the market-based strategy, and this conjecture is confirmed by Figure 3. While the chart for the macro-based strategy is essentially the same as in Figure 1, the market-based strategy no longer exhibits the big drop in the late 1990s and the zero total return in the second half of the sample. In general, both strategies are more aligned in the myopic case, indicating that the marked differences between the two strategies that are present in Figure 1 largely stem from different properties of the respective hedging demands. Therefore we next turn our focus to these hedging demands.

4.2 Determinants of the hedging demand

The hedging demand is determined by the dynamics of the investment opportunity set, i.e. by the covariance between asset returns and state variables. Our Markov-switching VAR model with parameter learning features two types of state variables.

First, for a given parametrization of the model, the Markov chain characterizes four possible states for the parameters A , μ and Ω in Equation(2). The state of the economy can switch in

every period and the covariance of regime switches and asset returns determines the hedging demand against regime switches. Second, the parametrization of the model changes over time because of the recursive updating of parameters. The investor in our model (suboptimally) disregards this source of variation in the investment opportunity set for tractability.

The hedging demand against regime switches depends on the dynamics of the Markov chain, which are characterized by two components: the return dynamics conditional on a given regime (i.e. the parameters A , μ , σ) and the probabilities of entering the respective regimes. From the perspective of a CRRA investor, these components capture two different aspects of the investment opportunity set: (a) how bad can the investment opportunities become (i.e. how do the possible states look like) and (b) how likely is a switch to one of these states? Given the complexity of the optimization problem in our model, we separate these two dimensions of the Markov chain by focusing on the expected returns μ and the transition probabilities P in the following. Figures 4, 5, 6, and 7 illustrate the differences in the hedging demand against regime switches for the macro-based and the market-based strategy.

Figures 4, 5, and 6 show properties of expected returns, i.e. of the estimated μ , for both strategies. For each asset class, we define the “dispersion in expected returns” as the difference between $\mu^{(j)}$ in the best and the worst possible state at each point in time. The “most extreme negative return” is defined as the worst possible expected return μ^j for each asset class and each point in time. Figure 4 depicts the time series of these two measures. Figures 5 and 6 show histograms resulting from these time series.

The differences between the macro-based and the market-based strategy are striking. In the market-based setup, the identified states are much more extreme and dispersed. Throughout the entire sample period and for every risky asset, there is a state in which these risky assets have large, negative expected returns. For the macro-based strategy, the expected returns in the worst possible state are only moderately negative most of the time (except for a few datapoints). Dacco and Satchell (1999) show that, in regime switching models, even a small number of wrong regime forecasts can lead to losing any advantage from a superior model. Figure 5 and 6 show that the market-based strategy exposes the investor to such a risk much more than the macro-based strategy.

For each point in time, Figure 7 shows the probability of leaving the current regime. To compute this probability, we define the current regime as the one with the highest probability according




Figure 4: Dispersion and most extreme negative returns

The top panel shows, at each point in time and for each asset class, the difference between the highest and the lowest mean return across all regimes. The bottom panel shows, at each point in time and for each asset class, the lowest possible mean return across all regimes. Both panels refer to the baseline case.




Figure 5: Histograms of dispersion of expected returns

The figure depicts histograms of the dispersion measure from Figure 4.




Figure 6: **Histograms of most extreme negative returns**

The figure depicts histograms of the measures of most extreme negative returns from Figure 4.




Figure 7: **Probability of leaving the current regime**

This figure shows, for each point in time, the estimated probability of leaving the current regime (defined as the regime with the highest probability at this point in time). The figure refers to the baseline case.

to the Markov chain estimation and then use the respective row in the transition probability matrix P_t . The transition probability matrix P_t is a key ingredient for the dynamics of the state variables π_t , as shown in Equation (6). Given the complexity of the portfolio optimization problem, which involves Bayesian learning and a solution based on Monte Carlo simulation, focusing on the transition probability matrix P_t provides the most efficient way of analyzing the dynamics of π_t . Figure 7 explicates that the probability of a regime switch between time t and $t + 1$ is generally higher with the market-based estimation as compared to the macro-based estimation, i.e. a change in the investment opportunity set is generally more likely for the former than for the latter.

Putting the two pieces together, we find that the investment opportunity set is much more volatile with the market-based strategy than with the macro-based strategy. A risk-averse investor, who wants to hedge against changes in the investment opportunity set, therefore invests less in risky assets, because the risky assets will perform particularly badly in the moment when the economy enters such a bad state. This effect is particularly severe after 2008 because the parameter learning (i.e., the quarterly re-estimation of parameters) changes the estimated parameters in particular in that year.

Figure 8 (top panel) shows the portfolio weights of the different asset classes for both strategies. It is clearly visible that, after the financial crisis in 2008, the market-based strategy invests mostly in the risk-free asset and in commodities, while the macro-based strategy tilts the portfolio to stocks and bonds again relatively quickly. The bottom panel shows the corresponding myopic strategies. Comparing the top and the bottom panel for the market-based strategy, we can see that the myopic portfolio comprises mostly risky assets, indicating that the overall risk-free portfolio allocation is due to a very strong negative hedging demand for stocks and bonds.

From the figure, we also see that the market-based strategy has a relatively large myopic demand for commodities, both before and after 2008. The fact that commodity returns were very low in the period after 2008 thus explains a large part of the overall bad performance of the market-based strategy. We interpret this finding as evidence that the regime identification based on stock and bond returns does not capture the dynamics of commodity returns well. The macro-based strategy does a lot better in this respect, it rightly puts zero weight on commodities (both myopic and hedging demand) after 2008. A robustness check with three asset classes only (i.e.




Figure 8: **Optimal portfolio weights**

The top panel shows the optimal portfolio weights for the macro-based and the market-based strategy in the baseline case. The bottom panel shows the respective portfolio weights for the myopic case.

without commodities), not shown here for brevity, confirms this result, as the difference between macro-based and market-based strategy is further (albeit not fully) reduced in this exercise.

Finally, we can draw a few more tentative conclusions from Figures 4 and 8. It seems that the effect of rare events like the 2008 financial crisis on the regime estimates is very long-lasting. This is in line with the general idea of costly parameter learning brought up by Collin-Dufresne et al. (2016). More recently, Kozlowski et al. (2020) have made a similar point, emphasizing the potential stickiness of belief changes after rare events. Interestingly, however, it seems that both the portfolio weights and the return parameter estimates stabilize more quickly under the macro-based estimation as compared to the market-based estimation, as more and more datapoints are included towards the end of the sample. This may indicate that the macro-based strategy requires less data to identify regimes reasonably well. This tentative conclusion is also confirmed by our robustness tests with different sample sizes, in particular shorter training periods, and with the pooled strategy, which will be discussed in Sections 5.1 and 5.3.

5 Robustness checks

Having explicated the main mechanism why the macro-based strategy outperforms the market-based strategy in our Markov-switching VAR setup with recursive updating of parameters, we now close the analysis with a set of robustness checks. We aim to document that the findings in

the previous two sections are not due to any esoteric choices that were made during the analysis, but are very general and indeed reflect deeply-rooted differences in macro-based as opposed to market-based regimes.

5.1 Pooled strategy

We start by introducing a third investment strategy that we label as “pooled strategy”. Until now, we have shown that, in a situation where the optimal solution cannot be obtained because of computational complexity, among the suboptimal strategies the macro-based strategy outperforms the market-based strategy. We have traced this outperformance to the key observation that the estimated Markov-switching VAR parameters for returns are more extreme when the regimes are identified with market data as compared to macro data. This raises the question whether the outperformance results from the regime identification being executed with *more* data or *different* data. For this question, it is important to keep in mind that we have deliberately not used any predictor variables in the VAR estimation conditional on the regime, so as to downplay as much as possible the impact of the predictability of returns by macroeconomic variables.

The question is not as trivial as it sounds. One has to bear in mind that all strategies that we discuss in this paper are suboptimal since we do not know the solution to the full optimization problem. We rank different, suboptimal strategies and among such suboptimal strategies more data does not always have to be better. It can also make the portfolio more exposed to the parameter updating problem or lead to more extreme regimes, exacerbating the hedging demand.


We therefore now analyze a strategy for which we pool macro and market data in the regime identification step. This strategy is equal to the other two strategies analyzed before, apart from one change: the regime identification (Step 1 of our iterative procedure) is done with macro and market data jointly, i.e., with four time series instead of two. The results are shown in Figure 9 and Table 3.

Most importantly, the pooled strategy does not outperform the macro-based strategy. It seems that the robustness of the macro-based strategy carries over to the pooled strategy, but the investor following the pooled strategy does not have superior information or a more stable estimation.

Table 3: **Robustness test: pooled strategy**

	Macro-based	Market-based	Pooled
Sharpe ratio	0.66	0.20	0.65
(p-value of the difference)	-	(0.04)	(0.96)
CEQ return	5.2%	-0.5%	5.1%
Total return	7.3%	1.7%	7.3%
VaR	15.5%	18.7%	15.8%
Turnover	0.62	0.33	0.47
Skewness	-1.24	-1.58	-1.19
Kurtosis	4.00	6.45	5.37

The table reports the performance of the different strategies in the baseline case and compares them with the pooled strategy. Sharpe ratio difference refers to the difference to the Sharpe Ratio of the macro-based strategy. The p-value of the Sharpe ratio difference is computed using the Ledoit and Wolf (2008) methodology. CEQ is the annualized certainty equivalent return, Total return shows the annualized return of the strategy over the whole period. VaR refers to the 5% value at risk and turnover captures the changes in the portfolio weights across different asset classes and over time.

Figure 9: **Robustness test: pooled strategy**

This figure shows the cumulative log returns of the different strategies in the baseline case with transaction cost and compares them with the pooled strategy.

The last point is also confirmed in Figure 10. The figure shows our measures of “dispersion” and “most extreme negative returns”, which are supposed to provide insights into properties of the estimated regimes. A comparison with Figure 4 suggests that the estimated regimes in the pooled strategy are very similar to the regimes in the macro-based strategy, when evaluated along these measures.

Stated differently, *replacing* market data by macro data in the regime identification step has the same effect on the portfolio performance as *adding* macro data to market data. Concerning our question above, we conclude that the outperformance rather results from using *different* data, not *more* data for the regime identification.




Figure 10: **Dispersion and most extreme negative returns**

The figure shows our measures of dispersion and most extreme negative returns, as introduced in Section 4.2, for the pooled strategy.

5.2 Different macro or market variables

In the baseline case, the macro-based strategy relies on consumption growth and inflation data, and the market-based strategy relies on stock and long-term corporate bond returns. But our main conclusions are robust to the use of alternative macro or market time series for the regime identification. Figure 11 and Table 4 show results for GDP growth instead of consumption growth and short-term Treasury bills instead of long-term corporate bond returns. The table and the figure show that our key conclusions are robust to these changes.

Interestingly, though, results for the stock-TBill strategy lie in between the macro-based strategy and the standard market-based strategy, for which the regimes are identified with stock and corporate bond returns. This is generally in line with the idea that the economic information contained in the time series of T-Bill rates is a mixture of the information contained in macroe-

Table 4: **Robustness test: different regime identifiers**

	Cons., inflation	Stock, bond	GDP, inflation	Stock, T-bill
Sharpe ratio	0.66	0.20	0.51	0.49
CEQ return	5.2%	-0.5%	3.1%	2.9%
Total return	7.3%	1.7%	6.5%	5.8%
VaR	15.5%	18.7%	20.7%	19.2%
Turnover	0.62	0.33	0.35	0.27
Skewness	-1.24	-1.58	-0.86	-0.59
Kurtosis	4.00	6.45	1.59	1.12

The table reports the performance of the different strategies in the baseline case using different variables for regime identification. The first two columns refer to the baseline strategies reported in Table 1. CEQ is the annualized certainty equivalent return, Total return shows the annualized return of the strategy over the whole period. VaR refers to the 5% value at risk and turnover captures the changes in the portfolio weights across different asset classes and over time.




Figure 11: **Robustness test: different regime identifiers**

This figure shows the cumulative log returns of the different strategies in the baseline case with transaction cost with different variables for regime identification.

conomic time series and in corporate bond returns. An analysis of the optimal portfolio weights and of the estimated regimes (similar to Figures 4-8, not shown here for brevity) confirms this interpretation.

5.3 Different sample periods

Another potential concern relates to the specific sample period used in the baseline case. To address this concern, we perform two robustness checks. First, instead of estimating the models with data starting in 1948, we shorten the sample for the estimation and start in 1970. Figure 12 and Table 5 show that our key results are robust to the shorter training period.

Second, we also check the portfolio performance measures on subsamples. Table 6 shows that our results are also robust to this exercise. For the case labeled 1995-2020 we use the period from 1948Q1 until 1994Q4 for the first datapoint, for the case labeled 2000-2020 we use the period from 1948Q1 until 1999Q4 for the first datapoint.

Table 5: **Robustness test: shorter training period**

	Macro-based	Market-based
Sharpe ratio	0.44	-0.05
(p-value of the difference)	-	(0.004)
CEQ return	2.0%	-6.8%
Total return	5.8%	-3.0%
VaR	22.5%	29.3%
Turnover	0.54	0.44
Skewness	-0.85	-3.76
Kurtosis	1.18	23.01

The table reports the performance of the different strategies in the baseline case, but with a shorter training period; 1970-1990 instead of 1948-1990. The p-value of the Sharpe ratio difference is computed using the Ledoit and Wolf (2008) methodology. CEQ is the annualized certainty equivalent return, Total return shows the annualized return of the strategy over the whole period. VaR refers to the 5% value at risk and turnover captures the changes in the portfolio weights across different asset classes and over time.




Figure 12: **Robustness test: shorter training sample**

This figure shows the cumulative log returns of the different strategies in the baseline case with transaction cost, but with a shorter training period; 1970-1990 instead of 1948-1990.

5.4 Transaction costs

We have already argued that transaction costs likely play a minor role for the outperformance of the macro-based strategy as compared to the market-based strategy, since the average portfolio turnover is in fact even higher for the macro-based strategy. Table 7 presents results for different values of the transaction cost parameter, where 0.50% refers to the baseline case analyzed in the previous sections.

The results confirm the intuition that we already derived from Table 1. The main results, i.e. the differences in Sharpe ratios and certainty equivalents, become even stronger if we disregard transaction costs, and they become weaker if we choose higher transaction costs. Actually, for transaction costs of 5.4%, the Sharpe ratios of the two strategies coincide. This reflects the higher average turnover for the macro-based strategy. In our specific setup with parameter updating, this implies that the estimation with macro data is less stable than with market data, requiring more extreme quarterly portfolio rebalancing on average. However, one can tentatively conclude from the performance depicted in Figure 1 that this parameter instability seems to improve the market timing ability of the investor following the macro-based strategy.

Table 6: **Robustness test: different out-of-sample periods**

	1995-2020		2000-2020	
	Macro-based	Market-based	Macro-based	Market-based
Sharpe ratio	0.66	0.17	0.45	0.09
(p-value of the difference)	-	(0.049)	-	(0.22)
CEQ return	5.3%	-1.1%	2.5%	-2.0%
Total return	7.8%	1.3%	5.0%	0.2%
VaR	16.6%	19.8%	18.0%	20.0%
Turnover	0.60	0.36	0.67	0.39
Skewness	-1.3	-1.74	-1.38	-2.07
Kurtosis	3.79	6.33	3.86	8.48

The table reports the performance of the different strategies in the baseline case, but with different out-of-sample periods. The p-value of the Sharpe ratio difference is computed using the Ledoit and Wolf (2008) methodology. CEQ is the annualized certainty equivalent return, Total return shows the annualized return of the strategy over the whole period. VaR refers to the 5% value at risk and turnover captures the changes in the portfolio weights across different asset classes and over time.

5.5 Autocorrelation in return dynamics

The baseline case allows for autocorrelation of returns in Step 2 (and 3) of our procedure (Equation (2)). This may be critical because of the concatenating of subsamples representing identical regimes, which is included in Step 1. As a robustness check, we therefore repeat our analysis and set $A = 0$ in Equation (2). The resulting portfolio performance is depicted in Figure 13.

As one can see from the figure, the performance is almost identical to the baseline case. Remember that all portfolio strategies discussed in our paper are globally suboptimal, as we have outlined in the introduction. In Step 3, the investor takes into account learning about the state of the economy, but disregards parameter learning. Therefore, one can also interpret Figure 13 as suggestive evidence that disregarding learning about autocorrelation parameters contributes only marginally to this suboptimality.

Table 7: **Robustness test: transaction cost**

Level of transaction cost	Sharpe ratio differences		Certainty equivalent returns		
	Macro – Market	Macro – Pooled	Macro	Market	Pooled
0.00%	0.51	0.03			
(p-value)	(0.03)	(0.86)	6.5%	0.2%	6.1%
0.25%	0.49	0.02			
(p-value)	(0.03)	(0.91)	5.8%	-0.1%	5.6%
0.50%	0.46	0.01			
(p-value)	(0.04)	(0.96)	5.2%	-0.5%	5.1%
1.00%	0.41	-0.02			
(p-value)	(0.06)	(0.93)	3.9%	-1.1%	4.1%
1.25 %	0.38	-0.03			
(p-value)	(0.08)	(0.88)	3.2%	-1.4%	3.6%
1.50%	0.36	-0.04			
(p-value)	(0.10)	(0.83)	2.6%	-1.8%	3.1%
5.40%	0.00	-0.22			
(p-value)	(0.93)	(0.44)	-8.3%	-7.1%	-5.2%

The table reports the performance of the different strategies in the baseline case with different levels of transaction cost. The p-value of the Sharpe ratio difference is computed using the Ledoit and Wolf (2008) methodology.




Figure 13: **Cumulative returns without autocorrelation in Equation 2**

This figure shows quarterly returns of the different strategies in the baseline case when the autocorrelation parameters A in Equation 2 are set to zero.

6 Conclusion

In their seminal contribution, Collin-Dufresne et al. (2016) show that parameter learning can be very costly: it generates long-lasting, quantitatively significant macroeconomic risks because it amplifies the impact of macroeconomic shocks on agents' marginal utility substantially. We document this key finding in a dynamic asset allocation setting where the parameter estimates for the data-generating process vary as more and more datapoints are observed and the sample size increases.

In such a setup with recursive updating of parameters, the globally optimal portfolio strategy cannot be determined due to the curse of dimensionality. Comparing different suboptimal investment strategies, we document that the portfolio performance can be improved substantially if the estimation relies on fundamental macroeconomic data instead of financial return data when determining the dynamics of regimes. Our findings show that, especially during highly uncertain times like the burst of the dotcom bubble or the 2008 financial crisis, the regime estimation based on financial market data lacks robustness and leads to a significant underperformance of the respective investment strategy.

Through a series of robustness checks, we show that using market data for the estimation amplifies the hedging demand with respect to regime changes, so that, in particular after the

disruptive financial crisis of 2008, the market-based portfolio strategy becomes very conservative, investing a large fraction of wealth in T-Bills. Overall, we conclude that, in a situation where the globally optimal solution to a portfolio problem is unknown because of the curse of dimensionality, macro data can enhance the risk-adjusted performance of the portfolio strategy.

References

- Amisano, G. and G. Fagan (2013). Money growth and inflation: A regime switching approach. *Journal of International Money and Finance* 33, 118–145.
- Ang, A. and G. Bekaert (2002). International asset allocation with regime shifts. *Review of Financial Studies* 15(4), 1137–1187.
- Ang, A. and G. Bekaert (2004). How regimes affect asset allocation. *Financial Analysts Journal* 60(2), 86–99.
- Ang, A. and M. Piazzesi (2003). A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics* 50(4), 745–787.
- Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance* 55(1), 225–264.
- Barsoum, F. and S. Stankiewicz (2013). Forecasting GDP Growth Using Mixed-Frequency Models With Switching Regimes. Technical Report 1.
- Bauer, M. and J. Hamilton (2017). Resolving the spanning puzzle. *Review of Finance* 21(2), 511–553.
- Bauer, M. and J. Hamilton (2018). Robust bond risk premia. *Review of Financial Studies* 31(2), 399–448.
- Beard, R. W., G. N. Saridis, and J. T. Wen (1998). Approximate solutions to the time-invariant hamilton-jacobi-bellman equation. *J. Optim. Theory Appl.* 96(3), 589–626.
- Bick, B., H. Kraft, and C. Munk (2013). Solving constrained consumption–investment problems by simulation of artificial market strategies. *Management Science* 59(2), 485–503.
- Bikbov, R. and M. Chernov (2010). No-arbitrage macroeconomic determinants of the yield curve. *Journal of Econometrics* 159, 166–182.

- Brennan, M. J. (1998). The role of learning in dynamic portfolio decisions. *Review of Finance* 1(3), 295–306.
- Brennan, M. J., E. S. Schwartz, and R. Lagnado (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control* 21(8-9), 1377–1403.
- Brennan, M. J. and Y. Xia (2001). Stock price volatility and equity premium. *Journal of Monetary Economics* 47(2), 249–283.
- Bulla, J., S. Mergner, I. Bulla, A. Sesboue, and C. Chesneau (2011). Markov-switching asset allocation: Do profitable strategies exist? *Journal of Asset Management* 12, 310–321.
- Chen, Z. and S. Jagannathan (2008). Generalized hamilton “cjacobi” cbellman formulation-based neural network control of affine nonlinear discrete-time systems. *IEEE Transactions on Neural Networks* 19(1), 90–106.
- Chernov, M. and P. Mueller (2012). The term structure of inflation expectations. *Journal of Financial Economics* 106(4), 367–394.
- Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review* 106(3).
- Coroneo, L., D. Giannone, and M. Modugno (2016). Unspanned macroeconomic factors in the yield curve. *Journal of Business & Economic Statistics* 34(3), 472–485.
- Dacco, R. and S. Satchell (1999). Why do regime-switching models forecast so badly? *Journal of Forecasting* 18(1), 1–16.
- de Pooter, M. D., F. Ravazzolo, and D. van Dijk (2007). Predicting the Term Structure of Interest Rates: Incorporating Parameter Uncertainty, Model Uncertainty and Macroeconomic Information. Technical report.
- Della Corte, P., L. Sarno, and I. Tsiakas (2010). Correlation timing in asset allocation with bayesian learning. Technical report, Technical report, Warwick Business School.
- DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal versus naive diversification: How inefficient is the $1/n$ portfolio strategy? *The review of Financial studies* 22(5), 1915–1953.
- Dergunov, I., C. Meinerding, and C. Schlag (2022). Extreme inflation and time-varying expected consumption growth. *forthcoming: Management Science*.

- Fallahi, F. (2011). Causal relationship between energy consumption (EC) and GDP: A Markov-switching (MS) causality. *Energy* 36(7), 4165–4170.
- Guidolin, M. and A. Timmermann (2005). Optimal portfolio choice under regime switching, skew and kurtosis preferences. *Working Paper*.
- Guidolin, M. and A. Timmermann (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control* 31, 3503–3544.
- Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics* 45, 39–70.
- Hess, M. (2006). Timing and diversification: A state-dependent asset allocation approach. *The European Journal of Finance* 12(3), 189–204.
- Honda, T. (2003). Optimal portfolio choice for unobservable and regime-switching mean returns. *Journal of Economic Dynamics and Control* 28, 45–78.
- Ilmanen, A., T. Maloney, and A. Ross (2014). Exploring macroeconomic sensitivities: How investments respond to different economic environments. *The Journal of Portfolio Management* 40, 87 – 99.
- Johannes, M., N. Polson, and J. Stroud (2009). Optimal filtering of jump diffusions: extracting latent states from asset prices. *Review of Financial Studies* 22(7), 2759–2799.
- Joslin, S., M. Pribsch, and K. J. Singleton (2014). Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks. *Journal of Finance* 69(3), 1197–1233.
- Kim, C.-J. and C. R. Nelson (1999). Has The U.S. Economy Become More Stable? A Bayesian Approach Based On A Markov-Switching Model Of The Business Cycle. *The Review of Economics and Statistics* 81(4), 608–616.
- Klein, M. W. and J. C. Shambaugh (2015). Rounding the Corners of the Policy Trilemma: Sources of Monetary Policy Autonomy. *American Economic Journal: Macroeconomics* 7(4), 33–66.
- Kole, E. and D. V. Dijk (2017). How to identify and forecast bull and bear markets? *Journal of Applied Econometrics* 32, 120–139.

- Kollar, M. and C. Schmieder (2019). Macro-based asset allocation: An empirical analysis. Technical report.
- Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020). Scarring body and mind: The long-term belief scarring effects of covid-19. *Jackson Hole Economic Symposium Proceedings*.
- Kritzman, M., S. Page, and D. Turkington (2012). Regime shifts: Implications for dynamic strategies (corrected). *Financial Analysts Journal* 68(3), 22–39.
- Ledoit, O. and M. Wolf (2008). Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance* 15, 850–859.
- Lewellen, J. and J. Shanken (2002). Learning, asset-pricing tests, and market efficiency. *The Journal of finance* 57(3), 1113–1145.
- Ludvigson, S. C. and S. Ng (2009). Macro Factors in Bond Risk Premia. *Review of Financial Studies* 22(12), 5027–5067.
- Martin, I. and S. Nagel (2022). Market efficiency in the age of big data. *forthcoming: Journal of Financial Economics*.
- Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3, 373–413.
- Munk, C. (2000, August). Optimal consumption/investment policies with undiversifiable income risk and liquidity constraints. *Journal of Economic Dynamics and Control* 24(9), 1315–1343.
- Munk, C. and C. Sorensen (2010). Dynamic asset allocation with stochastic income and interest rates. *Journal of Financial Economics* 96(3), 433–462.
- Nalewaik, J. J. (2015). Regime-Switching Models for Estimating Inflation Uncertainty. Finance and economics discussion series, Board of Governors of the Federal Reserve System (U.S.).
- Peijnenburg, K. (2018). Life-cycle asset allocation with ambiguity aversion and learning. *Journal of Financial and Quantitative Analysis* 53(5), 1963–1994.
- Pettenuzzo, D. and A. Timmermann (2011). Predictability of stock returns and asset allocation under structural breaks. *Journal of Econometrics* 164(1), 60–78.

- Song, D., H. Wang, and Z. Yang (2014). Learning, pricing, timing and hedging of the option to invest for perpetual cash flows with idiosyncratic risk. *Journal of Mathematical Economics* 51, 1–11.
- Tu, J. (2010, July). Is Regime Switching in Stock Returns Important in Portfolio Decisions? *Management Science* 56(7), 1198–1215.
- Turner, C. M., R. Startz, and C. R. Nelson (1989). A markov model of heteroskedasticity, risk, and learning in the stock market. *Journal of Financial Economics*, 3–22.
- Wang, N. (2009). Optimal consumption and asset allocation with unknown income growth. *Journal of Monetary Economics* 56(4), 524–534.
- Weitzman, M. L. (2007). Subjective Expectations and Asset-Return Puzzles. *American Economic Review* 97(4), 1102–1130.
- Xia, Y. (2001). Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation. *The Journal of Finance* 56(1), 205–246.
- Zeman, J. (2010). Estimating of bellman function via suboptimal strategies. In *2010 IEEE International Conference on Systems, Man and Cybernetics*, pp. 2406–2413.
- Zhou, H. (2020). Monetary Policy Autonomy: Identification and Cross-Country Evidence. working paper.