Research and Policy Notes 1/2021

Designing Macro-Financial Scenarios: The New CNB Framework and Satellite Models for Property Prices and Credit

Miroslav Plašil

Czech National Bank —— Research and Policy Notes —— 1/2021

CZECH NATIONAL BANK

www.cnb.cz

Reviewed by: Adam Geršl (Charles University, Prague)
Eva Hromádková (Czech National Bank)

Project Coordinator: Zuzana Gric

Designing Macro-Financial Scenarios: The New CNB Framework and Satellite Models for Property Prices and Credit

Miroslav Plašil *

Abstract

The paper sets out to present the Czech National Bank's new methodological framework for satellite models, i.e. models that link the macroeconomic scenario obtained from the core forecasting model with the evolution of key financial variables. Consistent macro-financial scenarios are particularly needed in macroprudential stress-testing. The paper describes the main underlying concepts of the new framework and provides further technical details on four newly deployed models for residential property prices and for bank loans in the main credit segments (housing loans, consumer loans and loans to non-financial corporations). The key advantage of the new approach is a shift to better-structured and more closely interrelated models. This should help maintain the internal consistency of the macro-financial scenario, facilitate communication of the assumptions behind the projections of financial variables and provide a high degree of robustness to structural changes in the economy.

Abstrakt

V článku je představen nový metodický rámec ČNB pro výstavbu tzv. satelitních modelů, které doplňují výstupy hlavního prognostického modelu ČNB o vývoj klíčových finančních proměnných. Toho lze využít zejména při přípravě konzistentních makroekonomických scénářů v zátěžovém testování. Článek popisuje základní principy nového přístupu a blíže dokumentuje čtyři nedávno nasazené satelitní modely pro projekce cen rezidenčních nemovitostí a bankovních úvěrů v třech hlavních úvěrových segmentech (úvěry na bydlení, spotřebitelské úvěry a úvěry nefinančním podnikům). Hlavní výhodou nového rámce je vyšší důraz na strukturální vlastnosti satelitních modelů a jejich hlubší vzájemná propojenost. Tyto charakteristiky by měly přispět k zachování vnitřní konzistence zátěžového scénáře, usnadnění komunikace předpokladů prognózy finančních proměnných a lepšímu vypořádání se strukturálními změnami v ekonomice.

JEL Codes: C51, C53, E37, E51.

Keywords: Gaussian process regression, macroprudential policy, satellite models, stress testing.

---

* Miroslav Plašil, Macroprudential Analyses Division, Czech National Bank (e-mail: miroslav.plasil@cnb.cz)
I am indebted to Eva Hromádková, Adam Geršl and colleagues from the Financial Stability Department for their comments and suggestions. I would also like to thank Michal Andrle for his collaboration on house price assessment methods. All remaining errors and omissions are my own.
The views expressed in this paper are those of the author and not necessarily those of the Czech National Bank.

1. Introduction

The majority of macroprudential authorities use a stress-testing framework as their main vehicle to assess system-wide resilience to a range of adverse economic shocks. The framework can also be employed to evaluate the financial health of financial institutions and other economic agents under the most probable ("baseline") scenario while supporting transparent policy-making.

Typically, the stress-testing exercise starts with the design of a coherent macroeconomic narrative, which is translated into actual numbers by the authority's core prediction model. However, the current generation of macro models rarely includes a financial sector and provides limited outcomes for relevant financial variables. To overcome this issue, stress tests commonly rely on a cohort of satellite (or auxiliary) models bridging macroeconomic and financial developments. The set of financial variables may include long-term (mortgage) interest rates, credit dynamics, asset prices and credit risk characteristics such as probability of default and loss given default.

This paper documents a new approach to building an internally coherent suite of satellite models used by the Czech National Bank (CNB) for its stress-testing framework. It introduces the general principles of the approach and provides some technical details on the satellite models for residential property prices and for the stock of credit in the main segments (housing loans, consumption loans and loans to non-financial corporations). Links to other satellite models in use or under development are also provided. Importantly, projections of property prices and loan dynamics serve as an important input into the stress-testing tools used to assess the resilience of households and non-financial corporations (Gregor and Hejlová, 2020; Siuda; 2020). These simulations can in turn generate trajectories of credit risk indicators such as probability of default.

The adoption of the new framework was motivated by the urge to give the satellite models a greater storytelling ability and higher internal consistency. A desire to obtain more reliable and economically sensible outcomes under plausible but severe scenarios was another trigger for redesigning the existing methodology. Satellite models are often estimated on samples with a paucity of financial crises, so their parameters carry only limited information about the behaviour of the economic system under severe financial distress. When the estimated model is later used to obtain the trajectories of financial variables under as yet unrealized extreme events, the outcomes require ex post expert interventions to bring them closer to economic intuition (in our experience, at least). In addition, the commonly used (regression-based) satellite models more often than not have difficulty capturing the non-linearities associated with financial developments through the cycle.

The new framework can be briefly characterized as an effort to move from a set of unrelated reduced-form models to a more structural and better-structured approach where the projections generated by the satellite models are interrelated. The link between house prices and mortgage loans can serve as a canonical example. These two quantities are clearly closely related through several channels, yet no formal link existed between them under the old approach. Their projections were based only on macro (non-financial) variables obtained from the core CNB macroeconomic model – a feature that all the former satellite models had in common.

Although the structure of the newly proposed models remains relatively trivial, it allows for capturing the non-linearities typical of the evolution of financial variables and offers enough flexibility to account for potential structural changes (e.g. lockdowns and repayment moratoria due to the Covid-19 pandemic). The rudimentary structure of the models also keeps their intellectual complexity at a reasonable level, a feature which should support overall economic intuition and easy interpretation of every aspect lying behind the final projection.

The shift towards fully structural satellite models is not complete, as some of the (sub)components are still left without deeper structural interpretation. We do not consider a semi-structural approach to be a drawback in our context. Although a more ambitious approach would be possible (at the cost of much greater complexity), semi-structural satellite models leave enough room for implementing expert judgement and incorporating extraneous information in areas where it can hardly be beaten by any model. There are also situations in which expert judgement and assumptions directly relate to the underlying story of the scenario, which is sometimes too subtle to be conditioned on or inferred from observed variables.[1] Where possible, we provide tools for probabilistic assessment of satellite projections that may help navigate expert judgement into plausible territory and guard against too extreme prior views. In addition, some robust empirical patterns are presented below to keep expert views within sensible corridors and to check for interrelations between the components modelled. These empirical regularities do not directly enter the satellite models as a quantitative input but have proved to be very useful for maintaining big picture consistency and safeguarding economic plausibility.

The remainder of the paper is structured as follows. Section 2 introduces the general principles of the new approach to building satellite models at the CNB and contrasts them with the older practice. Section 3 provides a broader picture of four recently deployed satellite models (residential property prices and the stock of loans broken down into the three main credit segments). It shows how structural relations are combined with time-series modelling techniques. The final section concludes. Technical details are relegated to Appendix A. A short illustration of how projections for residential property prices and the stock of credit were obtained during the first wave of the Covid-19 pandemic is provided in Appendix B.

2. The New Approach: Key Features and Related Literature

The CNB adopted a stress-testing framework for assessing the resilience of the Czech banking sector to system-wide shocks in 2003, following global trends in macroprudential oversight (Čihák and Heřmánek, 2005; Čihák et al., 2007). The stress-testing methodology has been gradually updated and extended for many additional features ever since (see Geršl and Seidler, 2010, and Geršl et al., 2012, and references therein). Along with improvements in methodology, the number of sectors and institutions subject to stress testing has increased immensely. At the time of writing, the CNB uses the stress-testing framework to assess resilience to risks in the banking sector, insurance corporations, pension funds, investment funds, households, non-financial corporations and government finance.

[1] For example, based on early information from market participants, the number of mortgage contracts was expected to fall dramatically during the lockdown in spring 2020 and to recover by the end of the year. This narrative simply could not be retrieved from the historical data, nor could it be fully derived from the macroeconomic scenario produced by the core model.

Coherent stress testing across a wide array of sectors requires a consistent and well-designed macro-financial scenario. The CNB moved away from historical ad hoc scenarios to fully fledged economic modelling in 2006 (Geršl et al., 2012) by supplementing the outcomes of the CNB core forecasting "g3+" model (Andrle et al., 2009; Brázdik et al., 2020) with projections of financial variables. The suite of satellite models that link a predefined macroeconomic scenario to financial developments has undergone several major re-specifications over the years, echoing a constant effort to improve the existing approach (Jakubík and Heřmánek, 2008; Geršl et al., 2012; Kučera et al., 2017; Panoš and Polák, 2019a,b). In addition to traditional satellite models, simulation tools for stress testing in the household sector and non-financial corporations' sector have been developed recently (Gregor and Hejlová, 2020; Siuda; 2020), providing useful information on the expected trajectory of credit risk variables and complementing the outcomes of the existing satellite models.

At present, satellite models are mainly used to produce trajectories for long-term yields (government bonds and interest rate swaps), residential property prices, credit growth and credit risk parameters (probability of default, PD, and loss given default, LGD) broken down by the main credit portfolios. The final scenario is then fed into the stress-test machinery to quantify the degree of resilience to selected financial risks.

Despite great efforts and an advanced econometric framework, some of the existing satellite models did not meet all the requirements of the stress-testing teams (the former practice is referred to as the "old approach" in the rest of the paper). The main shortcomings were the lack of a structural interpretation of the models, weak adaptability to the richness and subtleties of the narrative outlined by the scenario, a limited ability to take on board observed structural changes, unintuitive outcomes for severe economic downturns and problems with achieving internal consistency. In some cases, a lack of necessary detail was also a potential issue (e.g. the aggregate PD for the whole non-financial corporations sector vs. industry-level PDs).

One option was to re-estimate the existing reduced-form regressions, find a better set of explanatory variables or try to move to non-linear settings in order to obtain more intuitive outcomes. However, this would have hardly addressed all the concerns raised above. In particular, the former generation of mutually independent reduced-form models does not provide enough economic insights and remains too blunt to develop economic intuition. For these reasons, a new approach which takes a different path towards financial projections was adopted.

The new approach consists of a few general modelling principles which are adapted as much as possible to the individual satellite models. The rest of this section presents the key underlying ideas and contrasts them with the old approach. The next section then provides details on how these principles are reflected in four newly implemented satellite models. Strictly speaking, the methodological principles set out below only apply in full to the satellite models presented in this paper, but the effort to build a coherent framework is more general and penetrates into other auxiliary models as well. The satellite models under development also share the same philosophical starting point and (hopefully) all their future revamps will stick to the same modelling ethics, too.

A creative process, not a one-button-click procedure

Although quantitative methods still play a major role in the conception of the macro-financial scenario, expert judgement and the pursuit of a transparent economic narrative stand at the forefront of the scenario design under the new approach. Obtaining projections of financial variables from satellite models can be described as a quantitatively guided process where probabilistic and structural relations provide firm corridors for the underlying economic storyline. This "story-first" approach might be more time-consuming and require more thorough internal discussion than traditional model predictions, but it provides clear benefits.

Under the old approach, expert knowledge was also an integral part of the process, but in an unsystematic and less effective way. The trajectories obtained from satellite models almost always seemed implausible to the experts and often needed considerable manual ex post revisions. However, the experts' ability to rectify model outcomes "on the fly" proved to be very poor. Lacking quantitative rigour, economic intuition often turned out to be very wrong in explaining the economic rationale behind the manually adjusted projections and was simply not sufficient to get the projections generated by the models on an economically sensible path. Given this experience, expert judgement complemented with quantitative tools seems a more promising way to go than model projections obtained with one click of a button. In particular, the new approach should help us to exploit expert intuition in areas where it can potentially be very strong and to limit the cases where it may heavily fail.

Internal consistency with the CNB's core macro prediction ("g3+") model

In theory, the satellite models presented below can be linked to any macroeconomic forecast. However, internal consistency with the official macroeconomic prediction model is pursued in macroprudential stress-testing practice at the CNB. Projections of financial variables are made fully conditional on the forecast generated by the official macroeconomic model and therefore do not fall into the category of traditional forecasting (the goal of full consistency with the core macro-model prevails over real-world forecasting accuracy). In other words, obtaining accurate forecasts for financial variables when the underlying macroeconomic forecast has deviated considerably from reality should be considered a failure, not a success. Other models would probably have been used if the goal was to obtain the best forecasts possible.

These principles also relate to the important issue of macro-financial feedback loops and the need to model them in satellite models. Given the goal of consistency with the predefined macroeconomic forecast, feedback loop considerations clearly do not make much sense here, because of the implicit assumption that the underlying macroeconomic projection generated by the core model already includes them (even if it does not). The only role for satellite models in our paradigm is to recover the (most) likely trajectories of financial variables that would have been observed under the given macroeconomic scenario. One can think of satellite models as a specific sort of missing-data imputation method in an environment where only a subset of the variables is known to (or observed by) the analyst.

Internal consistency between satellite projections of financial variables

The new framework strives to come up with a system of interlinked satellite models where the trajectories of financial variables depend on both the macro projections and the outcomes produced by the rest of the satellite models. Internally consistent projections should strengthen the economic narrative of the scenario and provide formal modelling tools for narrowing the corridor for expert adjustments to the scenario (as they reduce the number of components which can be freely adjusted). This is in stark contrast to the older approach, which provided no guarantees of consistency among financial variables. Each satellite model generated fully independent projections based purely on macro variables and thus lacked important information from the other satellite models. The inconsistencies were potentially further aggravated by ad hoc expert interventions where manual adjustments to the outcomes of one satellite model were not translated into the other models. Under the new approach, any change made to one financial variable is automatically reflected in the remaining satellite models.

More structure

The catchphrase "more structure and less econometrics" perhaps best encapsulates the main paradigm shift between the old and new framework. The better-structured approach to satellite projections enables us to break down the whole system into smaller logical pieces ("modules") and find suitable representations for each of them. At the same time, it helps better describe the logical relations between the modules. For example, a projection for the stock of mortgage loans can be decomposed into a problem of finding a suitable (structural) model for newly provided flows of credit on the one hand and repayments of existing debt on the other. New credit volumes can be further factored as the average size of a mortgage and the number of new credit contracts and so on. The average mortgage size can then be functionally linked to house price developments, for instance.

The advantage of such an approach is threefold. First, better-structured models greatly enhance the accountability of the scenario and help better communicate its underlying assumptions. These assumptions can be changed for sensitivity analyses and counterfactual exercises, if desired. Second, it makes it easier to account for structural changes and allows one to implement the real (or expected) effects of economic policies. Lastly, the modular nature of the satellite models allows for future methodological refinements of individual parts of the system without the necessity to change its basic design and structure. Needless to say, a well-devised economic structure of satellite models leads to more cautious feature engineering and far superior longer-term forecasting properties than is the case for common satellite models.

Change in the preferred modelling choices

The satellite models still include modules without structural interpretation and need to be modelled in "regression-style" fashion. Our statistical approach, however, differs from the older practice in several aspects. First, all financial variables are directly modelled in levels, not in growth rates (regardless of their final presentation format). This supports the fact that levels (e.g. credit exposures), not growth rates, are the quantities of interest in stress testing. Although the issue is largely overlooked in applied research, growth rate regressions can often be problematic from an economic point of view (see, for example, Andrle and Plašil, 2019a, and Andrle and Brůha, 2017) and prevent practitioners from taking advantage of the huge benefits of the "trend-cycle" set-up. Among other things, such decomposition acknowledges the important economic wisdom that different frequencies are dominated by different economic phenomena.

The notion of frequencies also plays an implicit role in the estimation of model (hyper)parameters. It has been common practice to estimate the parameters of satellite models by minimizing the one-step-ahead prediction error. This might not, however, be optimal practice if frequencies are not equally important and model misspecification is a potential issue (see Haywood and Tunnicliffe-Wilson, 1997, and Andrle and Brůha, 2017). Given that stress tests draw on long-term projections with forecast horizons stretching up to five years, we make sure that the models are able to capture the mid-to-long-term behaviour of the underlying financial variables and downplay high-frequency dynamics (see below for how this is implemented in practice).

All reduced-form projections are newly guided by economically meaningful steady states or steady-state growth rates. This means that the projections take into account the current position in the cycle but gradually converge to its long-term (“trend”) dynamics. Steady states provide a transparent and easily communicated projection anchor that can be changed or updated if needed. Overall, empirical evidence shows that incorporating steady-state information into the estimation process can greatly enhance the forecasting performance of the model (Villani, 2009).

Our preferred method of inference in the area of satellite models is Gaussian process regression (GPR) – a Bayesian non-parametric regression tool popular in geostatistics (commonly known as *kriging*) and machine learning (Rasmussen and Williams, 2006). GPR offers great flexibility in modelling non-linear relationships, enables seamless implementation of various sorts of expert prior beliefs and provides straightforward tools for probabilistic assessment of the selected trajectories. Importantly, the relative likelihoods of different scenarios can easily be evaluated due to the computational tractability of the GPR model. The method also makes it possible to find the most probable scenario under the restriction that variable values in some periods were pre-set by experts. This can be helpful in situations in which experts have strong reasons to believe that short-term projections can be best captured by ad hoc considerations (see, for example, footnote 1).

3. Newly Implemented Satellite Models

This section provides background information on some of the newly deployed satellite models. Our ambition here is to set out their main structure and show potential avenues for further refinements. Although essential technical details on estimation and calibration are provided, the narrative mainly focuses on the underlying economic logic. Interested readers are encouraged to check the cited literature and technical Appendix A for details on the econometric and machine learning methods applied. A hands-on illustration of how projections can be obtained in practice is shown in Appendix B.

3.1 Residential Property Prices

The new satellite model for residential property prices is, in principle, a semi-structural trend-cycle model where the trend coincides with the fundamental value of residential property and the cycle can be linked to property overvaluation (undervaluation). While the fundamental value is modelled via a simple structural model, the model for overvaluation maintains its reduced-form nature.

We use the borrowing-capacity approach as our preferred concept for modelling fundamental developments in property prices, due to its simplicity, clarity and straightforward link to the official macroeconomic CNB forecast. This approach is used at the CNB as one of the methods to assess residential property prices and their overvaluation, but its area of application is much wider, including forecasting (for details and a justification of the approach, see Andrle and Plašil, 2019a, or Andrle and Plašil, 2019b; henceforth AP). Although the model does not rely on asset valuation theory and rather takes a macroprudential perspective on a property’s appraisal value, the results are quantitatively quite similar to those obtained by valuation methods.

The borrowing capacity approach is based on the premise that house prices reflect the capacity of credit-constrained households to borrow and pay for the property, where any increases in capacity are promptly priced in by the market. Although this approach might be heterodox from the viewpoint of valuation theory, surveys and anecdotal and empirical evidence suggest that this premise is actually very realistic in practice (AP, 2019b). The fundamental value of residential property is thus related to the maximum size of the loan that a prudent (median) household can safely borrow given their income, the prevailing market interest rate and prudent levels of borrower-based limits.[2] If observed prices are higher than borrowing capacity coupled with own funds, households are willing to accept higher financial risks on average, which may imply increased sensitivity to adverse shocks in the future. Overvaluation under this approach thus corresponds to a signal of systemic risk accumulation.

Determining the maximum attainable loan relies on the well-known mortgage contract calculations. A household can allocate a portion, α (the LSTI ratio), of its disposable income, Y_t, to service its monthly mortgage payment, A_t.

$$A_t = \alpha Y_t. \tag{1}$$

Given the monthly payment, A_t, the monthly mortgage interest rate, i_t, and the loan maturity (in months), N_t, the size of the attainable loan, L_t, is given by:

$$L_t = A_t \times \left[\frac{(1+i_t)^{N_t}-1}{i_t(1+i_t)^{N_t}}\right]. \tag{2}$$

The attainable loan, L_t, supplemented with own savings (the down payment) determines the amount of housing, V_t, a household can afford. For known values of the loan-to-value ratio, ϕ, attainable housing can always be expressed as:

$$V_t = \frac{1}{\phi_t} \times \alpha_t Y_t \times \left[\frac{(1+i_t)^{N_t}-1}{i_t(1+i_t)^{N_t}}\right] \equiv f(Y_t, i_t; \phi_t, \alpha_t, N_t). \tag{3}$$

The structural model (3) corresponding to fundamental developments in residential property prices (i.e. the trend) shows that prices in the long run are a function of only two variables (disposable income and mortgage rates) and three structural parameters (the loan-to-value ratio, the loan service-to-income ratio and the loan maturity).[3] Despite its simplicity, model (3) has important implications for residential property prices in the long run (see AP, 2019a,b). We briefly comment on those which are the most relevant for satellite models and house price projections.

- House prices will grow in line with growth in the nominal disposable income of households in the long run. With the loan-to-value ratio, the debt service-to-income ratio, loan maturity and mortgage rates all constant, residential property prices will generally grow faster than consumer prices, as their long-term anchor (i.e. steady-state growth) is determined by growth in nominal income.

[2] Note that these limits do not necessarily coincide with the binding (macroprudential) limits, as the latter typically constrain the upper tail of the distribution while usually leaving median household unaffected.

[3] It would be straightforward to incorporate the loan-to-income (LTI) ratio as well by defining the attainable price as the minimum of (3) and the price implied by the LTI ratio. However, for a plausible range of interest rates, equation (3) usually represents a more severe constraint.

- Fundamental house prices are non-linear in interest rates. The degree of non-linearity becomes particularly pronounced at low interest rates and the non-linear pattern is amplified by increasing loan maturity (see equation 3). This feature helps to explain why strong growth in property prices can be observed in a low interest rate environment even when a small change in mortgage rates occurs. It turns out to be essential for successful modelling of residential property prices under various scenarios.

- The flow of credit (newly granted mortgage loans reflecting current borrowing capacity), not the stock of credit, is the quantity closely related to house price developments. This provides a functional link between two different satellite models (or, more precisely, between two of their subcomponents) and shows that a more subtle connection between house prices and banks' total mortgage exposures is needed than common intuition would perhaps lead one to believe. The same conclusion has also been reached empirically (see Adalid and Falagiarda, 2018), confirming the theoretical implications of our simple structural model.

Practical implementation of model (3) proceeds as follows. We use the same data definitions and parameter settings as in AP (2019a), i.e. mortgage rates are set to the average interest rates for new mortgages and income is defined as disposable income per capita scaled by a constant factor of 1.65 to match the median income levels reported in mortgage loan applications. The prudent LTV and LSTI thresholds for the median Czech household are kept constant in the projections and are set to 0.35 and 0.8 respectively. The typical mortgage loan is assumed to have a maturity of 25 years or 300 months (see AP, 2019a, for a justification of the parameter settings).

Note that the choice of structural parameters determines the level of attainable prices but not their dynamics. This implies that once they are kept constant in the analysis, they only have an impact on the average (expected) size of overvaluation[4] without affecting its overall shape. If the expected value (i.e. the unconditional mean/steady state) of overvaluation is carefully tackled in the projection, the choice of structural parameters is not as crucial as it is for the purposes of overvaluation assessment. However, if structural changes are likely to occur, the parameters can be set differently in each period to fully benefit from the advantages of the structural model.

To obtain projections for attainable prices (equation 3) corresponding to the long-term trend in residential property prices, inputs for household disposable income and mortgage rates are needed. Projections of income can be directly obtained from the core CNB forecasting model, while projections of mortgage rates are based on the existing satellite model of long-term (15Y) government bond yields (Kučera et al., 2017). Government bond yields and mortgage rates have similar levels and exhibit strong comovement, although the relationship has weakened somewhat in recent years. In particular, the expected termination of the exchange rate commitment in 2017 triggered massive foreign investment in government bonds denominated in Czech koruna, which pushed their yields into negative or ultralow territory – considerably below mortgage rate levels. While expert judgement and market intelligence can be applied to short-term projections of

[4] Note that since we do not use a regression model, overvaluation is not forced to be zero on average. AP (2019a,b) show that the assumption of zero-sum residuals is actually quite unrealistic in this context.

mortgage rates, it is generally assumed that mortgage rates and long-term government bond yields will converge to similar levels again once the effects of koruna "overboughtness" fade away.

Comparing fundamental values with observed market prices determines the size of cyclical fluctuations in prices, or the "valuation gap". In line with common practice, the gap is expressed as a percentage of the fundamental value. To make the comparison possible, residential property prices need to be expressed in units of currency, not just as a dimensionless index number. Since absolute prices per square meter are only available for apartments in the Czech Republic, we had to adopt simplifying assumptions to implement the borrowing capacity approach for the whole residential property market. In particular, we rescaled the residential property price index (HPI) to an artificial price expressed in money terms by minimizing the difference between the size of overvaluation obtained for apartments (which is calculated on a regular basis) and that obtained for residential property as a whole.

Note that this one-off operation only rescales the level of the HPI without changing its dynamics. The overvaluation of the various residential property types (land, family houses and apartments) is not necessarily identical, but this is not a major issue here since the rationale behind the selected rescaling is to mimic a key economic feature of borrowing capacity. Namely, as observed in AP (2019a,b), the fundamental values determined by the borrowing capacity approach commonly form a lower bound for observed prices. This means that actual market prices are often overvalued but rarely, if ever, fall below the price implied by the borrowing capacity. The valuation gap tends to close fully over time, but the average size of overvaluation is not zero. This might be crucial for sensible elicitation of the steady state in time series models.

The valuation gap and its projections are modelled by Gaussian process regression (GPR – see technical Appendix A for a short, stand-alone exposition of GPR). It is a powerful non-parametric Bayesian technique applicable even in situations with very small data samples. In general, GPR solves the traditional regression problem:

$$y = f(x) + \epsilon, \tag{4}$$

where $f(x)$ is a latent (potentially highly non-linear) regression function describing the behaviour of the independent variable, y, based on the set of inputs, x. GPR captures the values of the regression function $f(x)$ at location x by a Gaussian process, which is a collection of random variables such that any finite subset of these variables follows a (multivariate) normal distribution. Due to the normality assumption, the Gaussian process is completely specified by its mean function $m(x)$ and covariance (kernel) function $K(x, x')$

$$f(x) \sim GP(m(x), K(x, x')). \tag{5}$$

Before observing the data, the mean function and covariance function of the Gaussian process are used to elicit priors about the nature of the regression function. The mean function can be used to elicit priors about the steady states of the process, while the covariance function defines the smoothness and periodicity (or dominant frequencies) of the process. After data arrival, predictions can be calculated based on the principles of conditional probability using closed-form formula for the conditional normal distribution (see Appendix A for details). The uncertainty intervals around the mean projections can serve as probabilistic corridors for potential expert interventions. The tractability of the multivariate normal distribution allows one to assign relative likelihoods of different expert-based scenarios vis-à-vis the mean projection. Moreover, if experts hold strong views or have relevant extraneous information on the size of overvaluation in some periods, it is straightforward to recalculate the mean projections and associated uncertainty conditional on this piece of evidence.

The most important part of the modelling process is to decide on a set of relevant explanatory variables, x, capturing swings in the risk perceptions of households that might be associated with overvaluation. This is a rather daunting task if one can only rely on the limited set of variables in the core CNB forecasting model or other satellite models. Out of the few possibilities available, we selected the unemployment gap as a rough proxy measure of households' perception stance. While this is only a rough approximation from the theoretical point of view, there is a strong empirical link between unemployment (the unemployment gap) and property prices (see, for example, Geerolf and Grjebine, 2014). An alternative measure capturing the difference between mortgage rates and short-term interest rates, which may approximate the ease of debt financing and households' willingness to buy a property, exhibits a similar pattern and thus provides quite similar information.

Besides the variables from the CNB core forecasting model, one can alternatively use time as a feature (explanatory variable) in GPR – converting the regression into a pure time series model which can be remotely likened to (atheoretical) autoregressive processes.[5] Although it is possible to work with all the variables in the GPR model at once, we do not do so and stick to a blend of the stacked regressions method (Breiman, 1996) and dynamic model averaging (DMA, Raftery et al., 2010). Stacked regressions represent an optimally chosen linear combination of simpler predictors, where the weights are determined by cross-validation. We generally follow this practice, with one small exception. We additionally use the forgetting factor known from the DMA and assign a lower (geometrically decaying) weight to the results of cross-validation on more distant training sets to account for the potentially slowly changing predictive capacity of the individual predictors across time or phases of the cycle. The one-year (four-periods) ahead squared prediction error is used as the metric to assess the size of the cross-validation error. This should capture medium-term developments in the valuation gap and suppress high-frequency noise.

Combining the long-term projection arising from the structural model (equation 3) with the reduced-form projection of the cycle (equations 4–5) it is possible to obtain the scenario-implied levels of house prices. When presenting the final outcomes, we do not allow for any manual interventions in the house price trend, but the projections of overvaluation can be subject to expert judgement if necessary. Among other things, this may also reflect strong preferences of the stress-testing team to achieve some targeted level of financial stress in the adverse scenario. However, GPR still provides probabilistic anchors that should keep ex-post interventions within plausible and consistency-preserving corridors.

3.2 Housing Loans

Housing loan dynamics are closely related to house price developments, but, as noted above, some care is needed to link these two quantities properly. In general, the concept of stock-flow reconciliation is essential to fully account for the fact that while new credit depends on the current economic conditions, loan repayment may depend on conditions and decisions related to the distant

[5] Note that careful elicitation of the steady state for the size of overvaluation is needed in this case.

past. This is particularly true for mortgage loans, where the length of a contract typically stretches over several decades.

The basic stock-flow reconciliation equation states that the difference in the stock of loans is equal to new loans, NL_t, minus loan repayments, LR_t, plus (or minus) other changes, OC_t.

$$\Delta L_t = NL_t - LR_t + OC_t. \tag{6}$$

Other changes may include instrument and sectoral reclassifications and balance-sheet write-offs, which can usually be assumed to be zero or negligible over the horizon of the projection. Using this slight simplification it is thus possible to obtain a scenario-consistent projection of the loan stock by separately modelling new loans and loan repayments. Separate modelling of the two quantities enhances economic intuition and provides greater modelling flexibility. For example, payment moratoria can easily be accounted for by making corresponding adjustments to the loan repayment schedule.

The volumes of new loans in a given period, NL_t, can be factored as

$$NL_t = avgNL_t \times nrNL_t, \tag{7}$$

where $avgNL_t$ stands for the average mortgage size at time t and $nrNL_t$ denotes the number of newly provided mortgage contracts. We relate developments in the average mortgage size to house price dynamics and assume that they both grow at the same rate in the projection.[6] Of course, in reality this assumption does not necessarily hold for at least two reasons. First, if the average LTV ratio changes, the link between prices and the size of a loan no longer holds exactly. Second, there might be situations in which households change their preferences and start buying smaller or larger apartments on average while the market price per square meter remains constant. Changes in the floor area of an average apartment can again erode the one-to-one mapping between the dynamics of house prices and loan size. Having said that, Figure 1 shows that the average mortgage size and residential property prices share similar long-run dynamics, confirming that the assumption is quite reasonable despite temporary deviations. We keep the LTV ratio and floor area constant in the projection, but adjustments can be made to these parameters if relevant information on future changes is available or when these changes are meant to be an integral part of the scenario tested.

[6] Information on the average mortgage size can be obtained from the regular CNB survey on retail loans secured by residential property or from Hypoindex (a local website that provides a basic set of statistics on mortgage loans). The information on the average mortgage size from the two sources matches quite closely. In practical implementation, we rely on data from Hypoindex due to their higher frequency and lower reporting lag.

Figure 1: Average Mortgage Size and House Price Index

Note: 2014 Q1 = 100
Source: Hypoindex, CZSO

The number of mortgage loans might be difficult to relate to the specific macro-financial variable, as it might – depending on many other factors – increase as well as decrease during economic slowdowns. For this reason, market intelligence (including expectations in the bank lending survey) and expert judgement play a crucial role in determining its future path, especially in the short run. In a way, the projection of the number of loans adds another assumption to the scenario and should be communicated as a part of its underlying story. In practice, we take into consideration several factors to navigate its projection into economically reasonable territory.

The first anchor relates to demographic trends and is given by the ratio of new mortgage contracts to the population of potential buyers (proxied by the size of the population aged 25–50 years).[7] With the exception of periods associated with excessive optimism and loose credit standards, this ratio is relatively stable over time (Figure 2). Although the short-term – and to some extent medium-term – projection of the number of contracts can be determined mainly by expert judgement, we let the projection converge to the long-term average of the ratio. In practice, we can again use the GPR method (with time as the explanatory variable), with the prior means set to the steady state described above. We also make sure that the long-term ratio does not lead to an explosive mortgage market penetration rate and converges to the penetration levels typically observed in advanced countries with similar mortgage market characteristics.

[7] Note that this assumes that demographic changes have a direct impact on the number of new mortgage loans in the long run. For the ratio to stay constant, the number of loans (the numerator) must adjust to the shift in the size of the population (the denominator). The results are not overly sensitive to other plausible definitions of the population of potential buyers.

Figure 2: Ratio of New Mortgage Contracts to Population of Potential Buyers

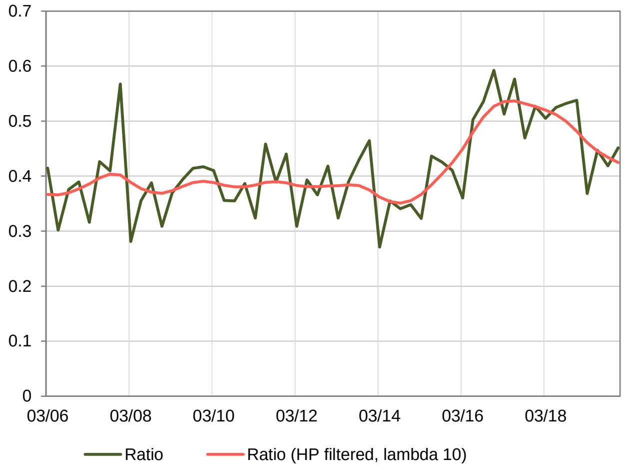

Note: Ratio of quarterly new mortgage contracts to the population of potential buyers in %.
Source: CNB, CZSO, author's calculations

Other factors taken into account include yields on alternative assets, inflation expectations, the correlation between the number of new mortgage loans and overvaluation and the strong empirical link between mortgage growth and unemployment (the unemployment gap). More specifically, in a low-interest rate environment accompanied by strong income growth and low yields on risk-free instruments, households can have strong incentives to invest in housing – combining available (low-interest bearing) savings with easily accessible financing. This behaviour may in turn lead to increased demand for residential property, supply-demand misalignments and eventually higher overvaluation. Even though the empirical correlation between overvaluation and the number of mortgage loans has not always been perfect, it has become somewhat stronger in recent years and may provide a useful guide for the projection.

Also note that the strong relation between household debt (mortgage debt in particular) and unemployment – which serves as another reasonable anchor – has not only been observed empirically (Figure 3) but is also supported by economic theory in models with endogenous money (Keen, 2016). This pattern should also be taken into consideration when deciding upon the projection of the number of new loans, as it clearly may narrow the range of economically meaningful alternatives. Besides systematic factors, the number of new loans can also depend on idiosyncratic shocks such as lockdowns of the economy due to the Covid-19 pandemic.

Figure 3: Unemployment Gap vs. Housing Loan Dynamics

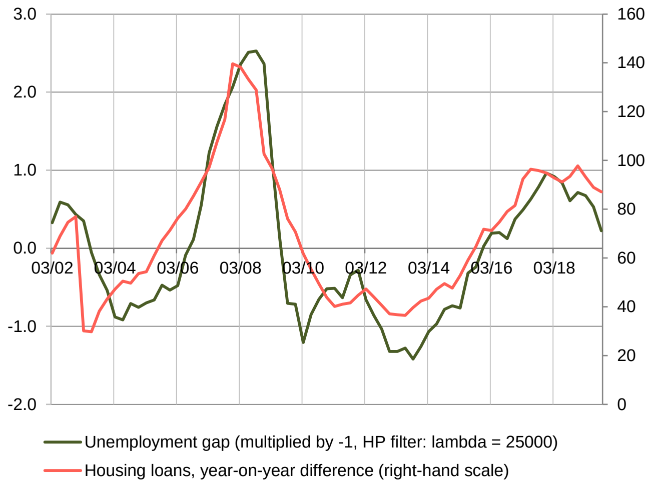

Note: Left-hand scale: %; right-hand scale: year-on-year change in CZK billions
Source: CNB, CZSO, author's calculations

Besides the projection of newly granted loans, the debt repayment schedule represents the second crucial ingredient of the credit stock dynamics in (6). To model the repayment process, we rely on the standard equations for regular loan payments, A_t, and the remaining principal, P_t, considering the whole stock of mortgage loans as a single mortgage.[8]

$$A_t = L_t \times \left[\frac{i_t^S(1+i_t^S)^{N_t^S}}{(1+i_t^S)^{N_t^S}-1}\right] \tag{8}$$

$$P_t = P_{t-1} - (A_t - P_{t-1} \times i_t^S) \tag{9}$$

The average stock interest rate, i_t^S, and the average maturity of the remaining stock, N_t^S, are used in (8) and (9). Information on average interest rates can be obtained from the banking statistics. The average stock maturity is generally not known but can be estimated using recent data and the stock-flow reconciliation equation (6). More specifically, it is straightforward to obtain the average maturity from (6), as it is the only unknown quantity in the equation once the data on credit stocks and new credit become available from the banking statistics. The estimates typically exhibit pronounced short-run oscillations around the long-term value, so in practice we use its smoothed version to derive the average maturity at the beginning of the projection.[9] Figure 4 shows that this approximation is still very good and allows for quite precise stock-flow reconciliation. The values of i_t^S and N_t^S can be held fixed in the projection, as the impact of their potential change is rather marginal. However, they can also be obtained recursively as weighted averages, where new loans (7) and remaining stock volumes (9) are used as weights. In the case of shocks to the repayment schedule, such as temporary payment moratoria, the schedule given by equations (8)–(9) can be

[8] Regular monthly payments are the overwhelming form of loan repayment in the Czech Republic, so this assumption should be a reasonable approximation. A more subtle approach would perhaps be needed in economies where this is not the case.

[9] This is mainly due to idiosyncratic effects (early repayments, credit defaults, other changes, different recording conventions for credit stocks and flows in the statistics etc.). Note that any systematic influence of these effects is reflected in (or captured by) the estimated level of the remaining average maturity.

flexibly adapted to better mimic reality. If the share of loans falling under payment moratoria is known, the schedule can be changed accordingly.

Figure 4: New Housing Loans and Estimate Thereof Based on Stock-flow Reconciliation with Estimated Average Maturity

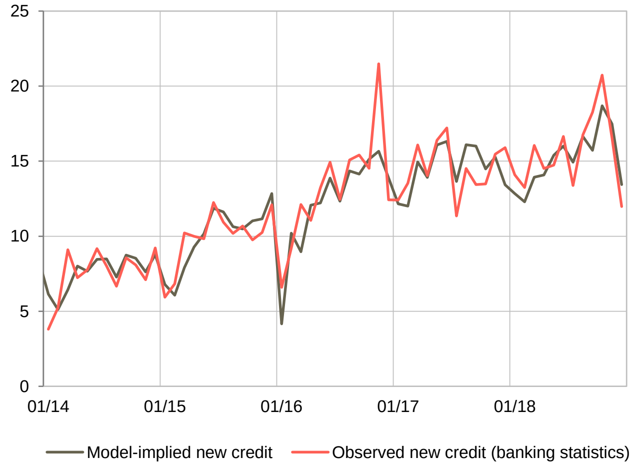

Note: Figures in CZK billions. New loans exclude refinanced and renegotiated loans. Stock-flow reconciliation is based on equation (6) using the estimated (smoothed) average maturity. Other changes are neglected.

Armed with projections of newly granted loans and loan repayments we can obtain projections of scenario-consistent housing loan exposures and their dynamics. At this point, we reiterate that given the structure of the model, it is easy to account for structural effects or breaks (see the examples above) as well as maintain consistency with other satellite models.

3.3 Consumer Loans

Similarly to housing loans, consumer loans are quite a homogeneous product with relatively low variation in their purpose, maturities and interest rates. The satellite model for consumer loans is thus built around similar principles and again makes use of equations (6)–(9). For this reason, only details on differences in the implementation for consumer loans are provided below.

Unlike for housing loans, the average size of a consumer loan is assumed to grow in line with CPI inflation and thus generally exhibits slower steady-state growth than housing loans do. The number of new loans is – analogously to housing loans – the most expert-guided quantity in the model. However, its overall path can be related to households' real consumption, unemployment and the interest rates on new consumer loans.[10] Again, we use Gaussian process regression to model the relation between the number of new contracts and the explanatory variables and apply the stacked regression technique to obtain a strong learner from weaker learners. The projection can be further anchored by the empirically observed strong positive correlation between new housing loans and new consumer loans. These two quantities rarely, if ever, diverge dramatically from a cloud of points determined by their strong pairwise correlation (Figure 5).

[10] Interest rates on new consumer loans are obtained as the 3M interbank rate (obtained from the core CNB forecasting model) plus a risk premium, which is modelled using a simple regression with variables taken from the core CNB model.

Figure 5: New Housing Loans vs New Consumer Loans: 2014 Q1–2020 Q2

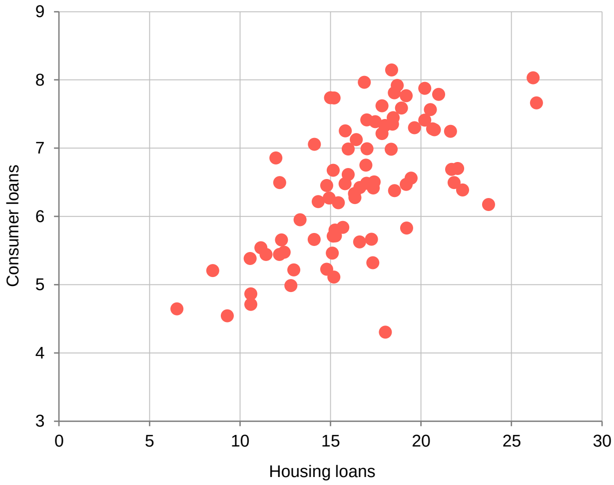

Note: CZK billions. New loans exclude refinanced and renegotiated loans.
Source: CNB

The repayment schedule again follows the same logic as before, but it has notably larger impact on the overall credit dynamics since the maturity of consumer loans is much lower than that of mortgages. It is also worth noting that information on the average maturity of the existing stock of consumer loans is available from commercial providers, which makes it possible to compare the accuracy of the estimation/approximation procedure described above with some hard data. The quantities are roughly equal, providing further support for the feasibility of the proposed estimation (approximation) procedure for the average maturity of the outstanding credit stock.

3.4 Loans to Non-financial Corporations

Loans to the non-financial corporations' sector can serve a wide variety of purposes, including investment, operational financing and mergers and acquisitions. This heterogeneity is reflected in considerable variation in maturities and interest rate levels across loan types, which in turn implies limited usefulness of their average values. Moreover, the link between flows and stock volumes is considerably looser than it is for household loans. This is due to how these quantities are recorded in the available statistics. In particular, data on newly granted loans cover information on all the new business of banks – a lot of which has no effect on the overall stock (e.g. credit lines). Moreover, misalignments in the time at which transactions are recorded in the stock and flow statistics can also be an issue. It is thus virtually impossible to apply the stock-flow reconciliation equation (6) or estimate a meaningful average maturity from the available statistical data.[11] For these reasons, we adopted a different approach to obtaining projections of corporate loans.

The underlying concept behind the satellite model for non-financial corporations' loans is that of the *credit impulse* (Biggs et al., 2010; Biggs and Mayer, 2013). It can easily be shown – both in theory and in practice – that once the information on credit is transformed into the credit impulse (essentially the scaled second difference of the credit stock), there is a remarkably strong relation

[11] The implied (remaining) average maturity of the stock would reach implausibly low levels of around one year. Available data from the credit register suggest that the maturity lies somewhere between four and five years.

between credit dynamics and growth in corporations' investment.[12] This is particularly true for medium and longer-term developments. Intuitively, if all investment (flow variable) is credit financed it must be directly linked to new borrowing, i.e. the flow of credit (Biggs et al., 2010). Against this background, growth in investment must then be closely related to growth in the flow of credit or the second difference of the credit stock. The relation would hold even if some part of investment was financed by own funds, as long as the financial leverage of firms is kept roughly constant. This is because investment and new borrowing would still grow at the same rate under these assumptions. The only thing that changes is the intercept in the "simple linear regression" model of the effect of the credit impulse on investment.

The pattern implied by the credit impulse proves to be empirically robust in the Czech context as well (Figure 6), although the impulse needs to be pre-filtered prior to the analysis in order to see this clearly in the quarterly data.[13] Because investment is a standard part of the forecast produced by the CNB core forecasting model, it seems straightforward to relate it to the credit impulse (via GPR) and retrieve a projection of the loan stock by back-transforming the predicted impulse. In other words, we first transform the series of credit granted to non-financial corporations into the credit impulse, regress it on investment growth obtained from the official CNB forecast and then transform it back to the stock of credit. No other variables (e.g. interest rates) are needed to run the regression, as these are already fully reflected in the scenario-predetermined forecast of investment. However, several additional remarks are still in order.

First, accounting for one-off events such as payment moratoria or a higher need for operational financing can be a little bit more challenging under the credit impulse approach than it was with the stock-flow reconciliation equation (6). But simple workarounds are still possible. Note that – other things being equal – these events directly increase the existing stock of loans and once they are quantified they have a directly measurable impact on the credit impulse.[14] The credit impulse associated with the projection of investment growth can then be adjusted ex post for the cumulative impact of one-off effects if these are assumed to be relevant over the projection horizon.

[12] This also implies that variation in credit is driven mainly by investment loans and investment-induced loans, while other types of loans adjust proportionally. Of course, this is not entirely true, but the approximation is still impressively good.

[13] In practice, we filter the credit impulse series using the HP filter with lambda set to 1 while penalizing excessive end-point variation by an additional constraint. As a check, we tried to reconstruct the stock of credit with a filtered and unfiltered credit impulse. There seems to be no systematic impact of the filtering operation on the overall stock levels.

[14] For example, changes in repayment volumes due to payment moratoria can be inferred from credit register data and available moratoria surveys. Information on higher financing needs can rely on data related to government support programmes and credit line drawdowns.

Figure 6: Investment of Non-financial Corporations and Credit Impulse

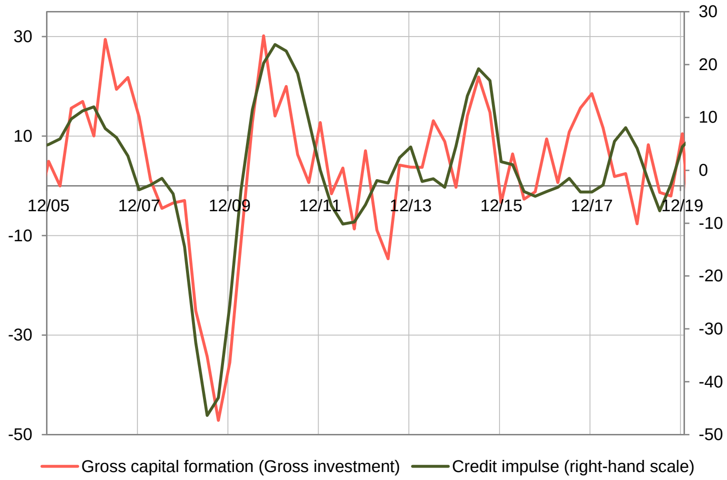

Note: Left-hand scale: investment, year-on-year growth in %; right-hand scale: credit impulse in CZK billions. The credit impulse is prefiltered by the HP filter with lambda set to 1.
Source: CNB, CZSO, author's calculations

The structural shift in the share of corporate loans denominated in foreign currencies and the increasing role of exchange rate fluctuations on the loan stock (when expressed in Czech koruna) are another issue that deserves some attention. This share has almost doubled in the last ten years and is currently hovering above 35%. This reflects an ongoing process of euroization in some parts of the economy, in particular in export-oriented industries. Loans in foreign currencies then serve as a simple form of hedging, with exporters striking a balance between two foreign currency flows (revenues and costs).

Taking into account the hedging motive, we assume that the shift to foreign currencies chiefly reflects a change in currency preferences but implies no change in the motivation for debt financing. This means that even though the growth rates of loans denominated in Czech koruna and foreign currencies might differ due to the increasing share of the latter, the overall stock is still driven by the same macroeconomic factors and can still be successfully retrieved from the credit impulse. In general, future changes in the exchange rate have relatively small effects on the overall stock if the exchange rate volatility is not dramatic, which allows us to abstract from exchange rate effects in most situations. However, if it is necessary to account for abrupt and very pronounced changes in the exchange rate, it is possible to incorporate them into the model provided the share of foreign currency loans in the overall stock is kept constant in the projection or set in line with expert views.[15]

Perhaps the most unsettling issue with the credit impulse approach is that the CNB core forecasting model only produces projections for the gross capital formation of the whole private non-financial sector (i.e. including housing investment of households). Its relation to non-financial corporations' credit impulse may be somewhat weaker (Figure 7). There are two possible ways to proceed: either one can work with the combined capital formation of households and non-financial corporations and accept a slightly worse regression fit (and higher uncertainty in the projection) or one can make expert interventions in the projection of investment, adjusting it for the capital formation of

[15] Several additional assumptions need to be made. In particular, one needs to assume the same repayment structure for loans in foreign currencies.

households. We stick to the first option for the time being, as the uncertainty related to the prediction of household fixed capital formation roughly offsets that associated with the worse regression fit.

Figure 7: Investment of Private Real-economy Sector and Credit Impulse

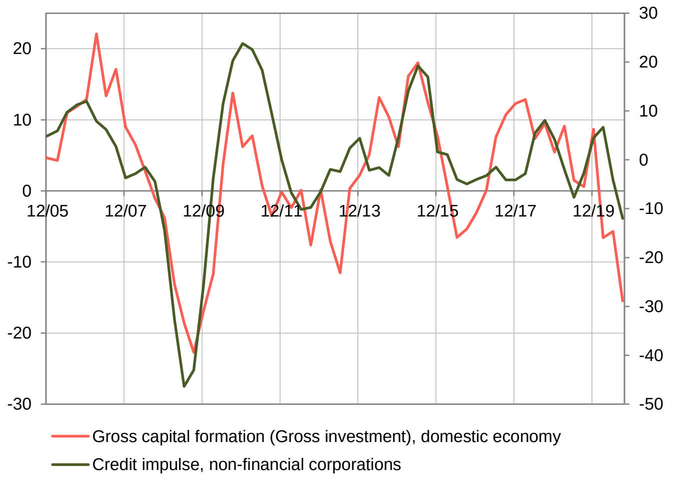

Note: Left-hand scale: investment, year-on-year growth in %; right-hand scale: credit impulse in CZK billions. The credit impulse is prefiltered by the HP filter with lambda set to 1.
Source: CNB, CZSO, author's calculations

Figure 8: Credit Growth versus Output Gap

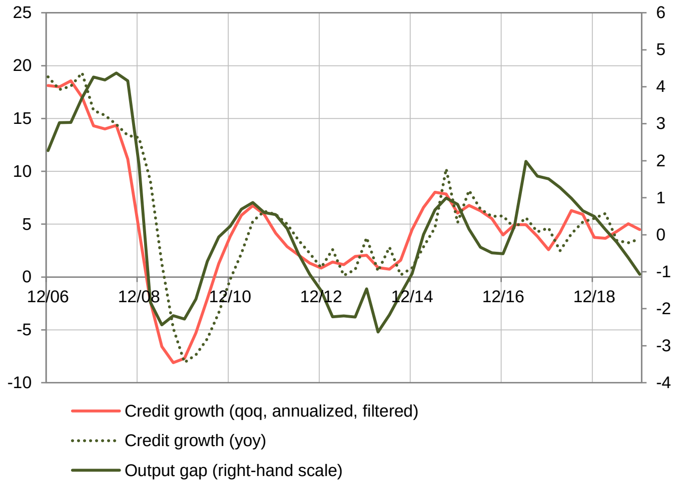

Note: Left-hand scale: credit growth in %; right-hand scale: output gap in % of potential output
Source: CNB, author's calculations

Again, the economic plausibility of the main projection can be assessed against other existing theory-grounded and empirically observed patterns. For instance, Biggs and Mayer (2013) show that under some assumptions growth in loans[16] can be related to the output gap. The CNB core forecasting model does not work with the concept of the output gap, but multiple internal estimates of its past and future values do exist (for example that based on the former semi-structural QPM model; see Beneš et al., 2003). Another option would be to take the GDP predictions from the core model and derive the gap by applying one's own preferred filtering technique, say a local linear

[16] More precisely, the output gap is related to growth in debt, so the empirical relation between the output gap and the change in bank loans may be blurred by developments in other forms of financing (e.g. debt securities issuance).

trend model with carefully calibrated steady-state values for potential growth. Figure 8 illustrates the relation between credit growth and the output gap obtained by the production function approach.

Last but not least, we also take some care with the long-term properties of the credit impulse-investment growth relation. Note that growth regressions can potentially imply quite divergent trajectories for the levels of variables in the long run even if the comovement in growth rates seems to be quite strong or even optimal under some statistical criterion. For this reason, we also check for the parametrized "level" version of the credit impulse model derived in Biggs et al. (2010; see eq. 4). Although their underlying structural model is a very simplistic two goods model, it allows the long-term behaviour of credit to be mimicked quite well (Figure 9). Again, one needs to bear in mind that only investment for the whole private non-financial sector is available from the core macroeconomic projection model. This may potentially worsen the model outcomes if the investment of households is remarkably strong, so some caution is in order. However, combining the information from the credit impulse with the long-term anchor still proves to be helpful most of the time. In practice, if the output obtained by regressing the credit impulse on investment growth is found to diverge significantly from the long-term level implied by the simple structural model, they can be reconciled by making small adjustments to the projected credit impulse.[17]

Figure 9: Long-term Relation between Credit and Investment in Simple Structural Model

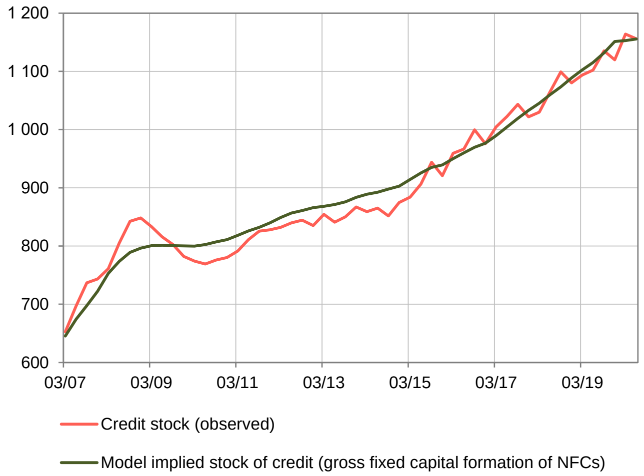

Note: CZK billions
Source: CNB, author's calculations

4. Conclusion

Stress tests have become a widely accepted workhorse for assessing the resilience of individual economic sectors to adverse system-wide economic shocks. To enhance the effectiveness of the exercise and facilitate communication of its outcomes to a wider audience, such assessment requires a well-defined and internally consistent macro-financial scenario. This paper presents the newly adopted framework of the CNB for designing a flexible, coherent set of financial variable projections built around the core macroeconomic "g3+" model. Following the framework, four

[17] This can be done formally by formulating a suitable minimization problem. In practice, however, it is done manually by trial and error.

recently deployed satellite models are described in greater detail. These include models for residential property prices and for banking loans broken down by the main credit segments.

The new approach offers greater flexibility in designing specific macroeconomic scenarios, facilitates the implementation of expert judgement in areas where it can be particularly strong and delivers more intuitive results under extreme but plausible events. By adopting a better-structured and interrelated set of satellite models, it is also possible to better account for structural changes in the financial system or in macroprudential policy conduct.

The four new satellite models presented in this paper exhibit both structural and reduced-form features and can thus be seen as semi-structural models. In general, the projection of financial variables is broken down into smaller parts (modules) which can be modelled individually and can be linked to each other through functional relations. Some of the modules can be redesigned later on without changing the main approach. Gaussian process regression is used to model the reduced-form modules. This allows for straightforward implementation of prior expert views and offers tools for flexible probabilistic assessment of various alternative projections. The paper also provides some robust empirical patterns that might be employed to keep projections within sensible corridors and to verify their consistency and overall plausibility. Although not explicitly used in the estimation process, the empirical patterns may be helpful in situations where expert judgement is used to partially override the results of statistical procedures.

The approach proposed in this paper can be adopted by macroprudential authorities or any other institution. The newly deployed satellite models can be linked to any macroeconomic forecast and can be adapted with respect to data availability in particular countries.

References

ADALID, R. AND M. FALAGIARDA (2018): "How Repayments Manipulate Our Perceptions about Loan Dynamics after a Boom". ECB Working Paper No. 2211/December 2018.

ANDRLE, M. AND J. BRŮHA (2017): "Forecasting and Policy Analysis with Trend-Cycle Bayesian VARs". Mimeo.

ANDRLE, M. AND M. PLAŠIL (2019a): "Assessing House Prices with Prudential and Valuation Measures". IMF Working Paper No. 19/59.

ANDRLE, M. AND M. PLAŠIL (2019b): "Assessing House Prices in Canada". IMF Working Paper No. 19/248.

ANDRLE, M., T. HLÉDIK, O. KAMENÍK, AND J. VLČEK (2009): "Implementing the New Structural Model of the Czech National Bank". CNB Working Paper 2/2009.

BENEŠ, J., T. HLÉDIK, D. VÁVRA, AND J. VLČEK (2003): "The Czech National Bank's Forecasting and Policy Analysis System". Chapter 4, pages 63–98, The Quarterly Projection Model and Its Properties.

BIGGS, M., T. MAYER, AND A. PICK (2010): "Credit and Economic Recovery: Demystifying Phoenix Miracles". Available at SSRN 1595980.

BIGGS, M. AND T. MAYER (2013): "Bring Credit Back into the Monetary Policy Framework!" *Political Economy of Financial Markets*, August 2013.

BRÁZDIK, B., T. HLÉDIK, Z. HUMPLOVÁ, I. MARTONOSI, K. MUSIL, J. RYŠÁNEK,T. ŠESTOŘÁD, J. TONNER, S. TVRZ, AND J. ŽÁČEK (2020): "The g3+ Model: An Upgrade of the Czech National Bank's Core Forecasting Framework". CNB Working Paper 7/2020.

BREIMAN, L. (1996): "Stacked Regressions". *Machine Learning* 24, pp. 49–64.

ČIHÁK, M. AND J. HEŘMÁNEK (2005): "Stress Testing the Czech Banking System: Where Are We? Where Are We Going?". CNB Research and Policy Note 2005/02.

ČIHÁK, M., J. HEŘMÁNEK, AND M. HLAVÁČEK (2007): "New Approaches to Stress Testing the Czech Banking Sector". *Czech Journal of Economics and Finance* 57, pp. 41–59.

GEEROLF, F. AND T. GRJEBINE (2014): "Assessing House Price Effects on Unemployment Dynamics". CEPII Working Paper No. 2014-24-December.

GERŠL, A., P. JAKUBÍK, T. KONEČNÝ, AND J. SEIDLER (2012): "Dynamic Stress Testing: The Framework for Testing Banking Sector Resilience Used by the Czech National Bank". CNB Working Paper 11/12.

GERŠL, A. AND J. SEIDLER (2010): "Stress Test Verification as Part of an Advanced Stress-Testing Framework." In: Czech National Bank, Financial Stability Report 2009/2010: pp. 92–101.

GREGOR, J. AND H. HEJLOVÁ (2020): "The Household Stress-test". Thematic Article on Financial Stability 4/2020, CNB.

HAYWOOD, J. AND G. TUNNICLIFFE-WILSON (1997): "Fitting Time Series Models by Minimizing Multistep-Ahead Errors: A Frequency Domain Approach". *Journal of the Royal Statistical Society*, Series B (Methodological) 59(1), pp. 237–254.

JAKUBÍK, P. AND J. HEŘMÁNEK (2008): “Stress Testing of the Czech Banking Sector”. *Prague Economic Papers* (3), pp. 195–212.

KEEN, S. (2016): “Modeling Financial Instability”. In: Malliaris, A. G., Shaw, L. and H. Shefrin (editors): The Global Financial Crisis and Its Aftermath: Hidden Factors in the Meltdown, Oxford University Press.

KUČERA, A., M. DVOŘÁK, L. KOMÁREK, AND Z. KOMÁRKOVÁ (2017): “Longer-term Yield Decomposition: An Analysis of the Czech Government Yield Curve”. CNB Working Paper 12/2017.

PANOŠ, J. AND P. POLÁK (2019a): “How to Improve the Model Selection Procedure within a Stress Testing Framework”. CNB Working Paper 9/2019.

PANOŠ, J. AND P. POLÁK (2019b): “The Impact of Expectations on IFRS 9 Loan Loss Provisions”. CNB Research and Policy Note 3/2019.

RAFTERY, A. E., M. KÁRNÝ, AND P. ETTLER (2010): “Online Prediction Under Model Uncertainty via Dynamic Model Averaging: Application to a Cold Rolling Mill”. *Technometrics* 52, pp. 52–66.

RASMUSSEN, C. E. AND C. K. I. WILLIAMS (2006): “Gaussian Processes for Machine Learning”. The MIT Press.

SIUDA, V. (2020): “A Top-down Stress-testing Framework for the Nonfinancial Corporate Sector”. CNB Working Paper 12/2020.

SCHULZ, E., M. SPEEKENBRINK, AND A. KRAUSE (2018): “A Tutorial on Gaussian Process Regression: Modelling, Exploring, and Exploiting Functions”. *Journal of Mathematical Psychology* 85, pp. 1–16.

VILLANI, M. (2009): “Steady-state Priors for Vector Autoregressions”. *Journal of Applied Econometrics* 24(4), pp. 630–650.

Appendix A

A.1 Gaussian Process Regression (GPR)

We consider a standard regression set-up where the dependent variable (output), y, is related to a set of explanatory variable(s), x, through a latent (unknown) regression function, $f(x)$:

$$y = f(x) + \epsilon, \tag{A1}$$

where $\epsilon \sim N(0, \sigma_\epsilon^2)$ is an i.i.d. error term independent of $f(x)$.[18] In Gaussian process regression, it is assumed that the regression function $f(x)$ follows a Gaussian process (Rasmussen and Williams, 2006; Schulz et al., 2018). This means than any two or more values of $f(x)$ at locations x follow a multivariate normal distribution. The Gaussian process is fully defined by the mean and covariance function

$$GP(m(x), K(x, x')). \tag{A2}$$

The mean function $m(x)$ captures the expected value of the function at location x

$$m(x) = \mathrm{E}[f(x)] \,. \tag{A3}$$

while the covariance function (or kernel of the Gaussian process) reflects the dependence between the function values at different locations of x.

$$K(x, x') = \mathrm{E}\big[\big(f(x) - m(x)\big) - \big(f(x') - m(x')\big)\big]. \tag{A4}$$

The choice of a suitable kernel and mean function before seeing the data determines the prior views about the smoothness, dominant oscillations, periodicity and other features of the function that the expert expects to observe in the data.

Once data become available, one can use the Gaussian process to make predictions. Under the assumption of independence of the error term and the regression function, the mean of y is $m(x)$ and the variance is equal to

$$Var(y) \equiv K_y = K(x, x') + \sigma_\epsilon^2 I. \tag{A5}$$

The joint distribution of the observed output values y and the predicted values f^* at the chosen locations x^* follows a multivariate normal distribution by definition of the Gaussian process. Following the properties of the multivariate normal distribution, the joint distribution can be partitioned as follows

$$\begin{bmatrix} y \\ \mathrm{f}^* \end{bmatrix} \sim N\left(\begin{bmatrix} m(x) \\ m(x^*) \end{bmatrix}, \begin{bmatrix} K(x, x') + \sigma_\epsilon^2 I & K(x, x^*) \\ K(x^*, x) & K(x^*, x^{*\prime}) \end{bmatrix}\right). \tag{A6}$$

Making predictions and quantifying the associated uncertainty then entails deriving the conditional distribution $\mathrm{f}^* | y, x, x^*$, which is again multivariate normal and is available in closed form:

[18] The implementation of different assumptions about the error term poses no difficulties in the GPR framework.

$$\mathrm{f}^*|y,x,x^* \sim N\left(\overline{\mathrm{f}^*}, Var(\mathrm{f}^*)\right)$$

$$\overline{\mathrm{f}^*} = m(x^*) + K(x^*,x){K_y}^{-1}\left(y - m(x)\right) \tag{A7}$$

$$Var(\mathrm{f}^*) = K(x^*,x^{*\prime}) - K(x^*,x){K_y}^{-1}K(x,x^*).$$

The functional form of the kernel, $K(x,x')$, should be chosen with respect to the regression problem at hand. The common assumption is that the correlation between two values of the regression function decreases with the distance between them. In other words, points lying closer to each other are expected to take more similar values than those further away from each other. In practice, several default kernel classes are used, based on prior views about the underlying properties of the regression function. The specific behaviour of the function can be tuned by a set of hyperparameters. In machine learning, the radial basis kernel ranks among the most popular class of kernels for modelling smooth stationary functions and is also the kernel used throughout this paper:

$$K(x,x') = \sigma_f^2 \exp\left(-\frac{\|x-x'\|^2}{2\lambda^2}\right). \tag{A8}$$

Two hyperparameters of the kernel (A8) – σ_f^2 (the variance of the signal) and λ^2 (the length-scale) – determine the smoothness and overall variability of the regression function. Their values can be either set in line with prior views or inferred from the available data. The estimation can be carried out by using full Bayesian inference or cross-validation. However, estimates are most frequently obtained by maximizing the marginal (log-)likelihood (together with estimating the unknown error variance, σ_ϵ^2):

$$\log p(y|x,\sigma_f^2,\lambda^2,\sigma_\epsilon^2) = -\frac{1}{2}y'K_y^{-1}y - \frac{1}{2}\log\left|K_y^{-1}\right| - \frac{n}{2}\log 2\pi. \tag{A9}$$

This is also the approach used in this paper, but cross-validation provided broadly similar results. The first term in (A9) captures the goodness-of-fit of the model, while the second penalizes model complexity and serves as a regularizer.

A.2 Stacked Regressions

Although GPR can easily handle a higher number of explanatory variables, we pursue a rather different approach here. This is because we want to account for the potentially time-varying explanatory power of individual inputs across individual stages of the cycle. For this reason we build a collection of weaker learners consisting of models with one input (the explanatory variable) only and obtain a stronger model by the method of stacked regressions (Breiman, 1996).

Stacked regressions resemble model-averaging techniques where the final prediction is obtained as a weighted average of the outcomes generated by the individual models. However, the existence of a single "correct" model associated with the data-generating process is not assumed here, and the collection of (non-negative) weights serves only as a tool for combining weaker learners. The weights assigned to the individual models are obtained by time series cross-validation, which is based on the four-step-ahead prediction error. Typically, the cross-validation error is calculated by taking the simple average of the prediction errors. We depart from this approach and assume that recent prediction errors are more relevant than those observed in the past. This reflects a belief that the predictive performance of the models is likely to vary across the business or financial cycle. To implement this idea, we draw on the prediction and updating equations known from Bayesian model averaging (Raftery et al., 2010). In particular, we slightly modify the updating equation for the weight assigned to model k at time t, $\pi_{t|t,k}$, and update it using the expression

$$\pi_{t|t,k} = \frac{\pi_{t-1|t,k} \times p_k(y_t|y^{t-4})}{\sum_{l=1}^{K} \pi_{t-1|t,l} \times p_l(y_t|y^{t-4})}. \qquad \text{(A10)}$$

where $p_k(y_t|y^{t-4})$ is the four-step-ahead predictive normal density for model k evaluated at y_t. The predictive density obtained from Gaussian process regression is normal by construction and is easy to evaluate. The prediction equation for the weights remains the same as in dynamic model averaging (see Raftery et al., 2010, for details).

Appendix B

This part illustrates how projections of financial variables within an expert-guided framework can be obtained in practice. In order to demonstrate the importance of extraneous information and the role of expert judgement, we describe the construction of projections of financial variables in the baseline macroeconomic scenario produced at the time of the outbreak of the Covid-19 pandemic in spring 2020. Macroeconomic figures for 2020 Q1 were not yet available at that time, but mushrooming lockdowns across the globe and deteriorating values of many leading indicators were already heralding an unprecedented economic slowdown. Although general expectations at that time were suggesting a relatively a fast recovery after a considerable but temporary drop in economic activity, macroeconomic forecasts were marked by huge uncertainty. Under these circumstances, it would have been unwise to rely fully on historical patterns and disregard other information or expert knowledge. The modelling process presented below (in a streamlined fashion) closely follows the real-time design of the projections but sometimes diverts from it in order to demonstrate some of the quantitative tools available to experts to guide their decisions. Due to space constraints, we limit our attention to the models for property prices, housing loans and loans to non-financial corporations. The satellite model for consumer loans shares a lot of similarities with that for housing loans and is therefore left out.

B.1 Property Prices

As explained in Section 3.1, the trend in property prices is given by the simple structural model (3), while the cycle (or the "valuation" gap) follows a Gaussian process (5). If the structural parameters in (3) are kept constant (as in our case), the trend can be readily obtained by inserting the forecasts of household income and mortgage rates into (3). The income of households came from the official CNB forecast produced in spring 2020. The forecast was prolonged by one year (a third year) by letting the forecast converge to the 5% steady state adopted by the CNB in its core prediction model. Since the forecast already assumed temporarily negative year-on-year income growth in 2020 Q2 due to deteriorating economic conditions, the income levels were smoothed prior to their insertion into (3). This reflected the assumption that a short-lived slump in household income should not (have time to) exert fundamental pressures on borrowing capacity and attainable housing. The assumption also implies smoothness in the trend component.

Mortgage rates were derived from the projections of 15Y government bond yields, which historically follow the same path. These projections come from a closely related satellite model (Kučera et al., 2017) which draws on the official CNB forecast of short-term interbank rates and other macroeconomic variables. It was assumed that the temporary gap between mortgage rates and bond yields caused by "overboughtness" of Czech koruna (see Section 3.1) would eventually close and both quantities would attain the same values at the three-year horizon. This implied a continuing decline in mortgage rates throughout 2020, stagnation in the first half of 2021 and a gradual increase at the far end of the forecast horizon. Using structural model (3) and the forecasts of mortgage rates and household income, it is easy to verify that the official macroeconomic CNB forecast was consistent with a steady upward trend in fundamental property values despite the expected economic slowdown (Figure B1).

By contrast, the size of overvaluation (the cycle) was generally expected to decrease somewhat due to a drop in mortgage contracts and lower demand amid uncertain economic prospects. Still, initial perceptions of the Covid-19 shock did not suggest long-term damage to the housing market, so it was quite reasonable to assume that renewed demand pressures would drive the overvaluation upwards again once the initial effects of the shock subsided (most notably in an environment of persistently low supply, which was about to be further eroded by lockdown barriers to property construction). Taking this into consideration, the steady-state size of overvaluation in the Gaussian process regression (GPR) with time as the explanatory variable was set to 10%. This corresponds to the level observed prior to the Covid-19 shock. The above narrative was also roughly in line with the course of the GPR projection based on the unemployment gap (Section 3.1), although the latter implies higher overvaluation at the far end of the projection.[19] Stacking the two GPR regressions using the weights (A10) presented in Appendix A, one obtains the final model-based projection of the valuation gap. This projection was later partially modified by expert judgement in order to incorporate a belief that the overvaluation would drop slightly more in 2020, but overall the changes introduced by experts were negligible.

Figure B1: Projection of Fundamental and Observed Property Prices

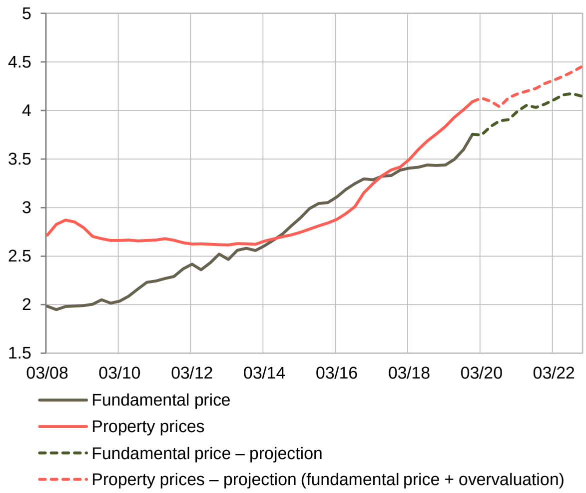

Note: Artificial price in CZK billions

With the benefit of hindsight, it is apparent that the baseline scenario undershot the actual surge in property prices by a large margin. However, a few comments are in order. First, the final projection (Figure B1) still implied positive year-on-year growth in property prices over the entire (three-year) forecast horizon and a considerably higher property price level at its far end. This is something many other models failed to deliver at that time, as they largely predicted property prices to fall (or to remain flat at best). Second, it is important to stress again that the projection of property prices does not represent a stand-alone forecast and should always remain contingent on the macroeconomic scenario and its narrative. A closer look at the fundamentals shows that the actual disposable income of households was higher and mortgage rates were lower than forecasted. Model (3) with observed rather than forecasted quantities would alone have explained the relatively

[19] This result should be viewed with caution due to general difficulties with estimating trends and gaps in an environment of massive shocks such as the Covid-19 pandemic.

fast growth in 2020. Moreover, the storyline of the baseline scenario was certainly more consistent, with some drop in overvaluation rather than a further increase.[20]

It might be useful to inspect some measures of uncertainty related to the projection and show how they might guide experts in their decisions to override the model-based forecast. Traditional (i.e. point-wise) 90% credible intervals are depicted in Figure B2 (left panel). One can see that the regression output is associated with relatively large uncertainty, in particular when time is used as the explanatory variable and no anchor from the official CNB macroeconomic forecast is provided. However, the point-wise credible intervals do not tell the full story. Note that the whole joint predictive distribution is known in the case of GPR (not just a point-wise distribution). This allows us to compare the plausibility of the projections based on their relative likelihood. In other words, not all the values and their combinations within the credible intervals are equally likely. Figure B2 (middle panel) depicts the relative likelihood for three different alternative scenarios vis-à-vis the mean of the Gaussian process. While all the projections take values within the 90% credible intervals, their economic plausibility given the model's predictive distribution is dramatically different. Another important observation is that the width of the credible interval crucially depends on prior beliefs. If experts truly believe that the valuation gap returns to its steady state at the end of the forecast horizon (or in any other arbitrary quarter in the future), it is possible to treat the steady-state value as a fictitious (dummy) observation. This naturally reduces the uncertainty, since there is now a fixed point where all the projections have to end (Figure B2, right panel).

Figure B2: Credible Intervals for Valuation Gap and Related Measures of Uncertainty

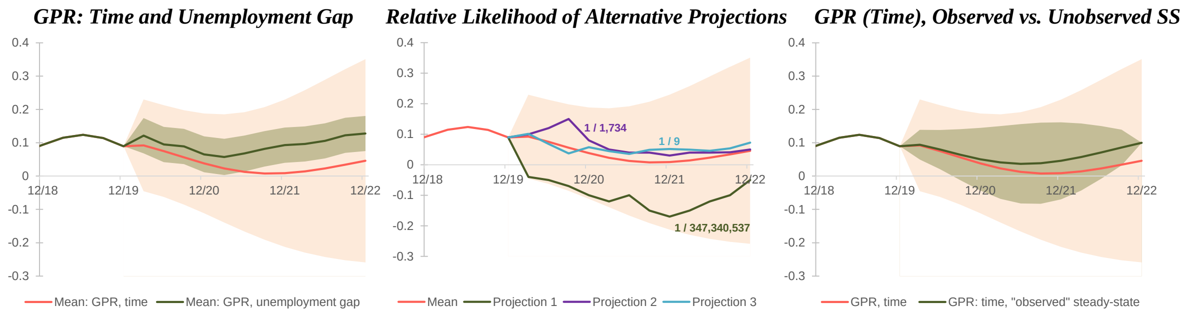

Note: Projections of the "valuation" gap based on Gaussian process regression (GPR). Positive values of the gap imply overvaluation of property prices. The gap is presented in %/100 (i.e. 0.1 = 10%). SS stands for steady state. Left panel: GPRs featuring time and the unemployment gap as independent variables. The other panels feature only time as the independent variable.

B.2 Housing Loans

The satellite model of housing credit is based on two building blocks – new credit and debt repayments. New credit volumes can be further factored as the average mortgage size and the number of new contracts (Section 3.1). We imposed one-to-one mapping between growth in property prices and the average mortgage size while making specific assumptions about the number of new contracts. The first-hand experience of banks at the start of pandemic pointed to a dramatic

[20] For example, the inflation forecast produced in spring 2020 was well below the current levels, implying considerably higher (non-negative) real mortgage interest rates and lower motivation of households to search for inflation-protected investments. The same applies to the unemployment rate, which decreased considerably less than predicted.

drop in mortgage contracts due to general uncertainty and the physical shutdown of banking outlets, which limited access to financial services. The first evidence-based estimates by bank officials concurred on a potential temporary decline of up to 50–70% in 2020 Q2 followed by normalization of the situation by the end of 2020.[21] Given the specific nature of the Covid-19 shock and the fact that no other reliable information was available, a temporary decline of 50% in 2020 Q2 followed by fast rebound in the subsequent quarters was expertly chosen as the most probable scenario. The steady state was set ad hoc slightly above the historical average to compensate for the temporary shortfall in new contracts (while assuming a fixed size of the population of potential buyers). Figure B3 (left panel) reveals that the pure model-based projections and the projections based on extraneous information backed by expert judgement differ dramatically in this case. The historical patterns alone could not account for the unprecedented shock. Also, note how the credible intervals are flexibly recalculated within the GPR framework once the extraneous information of a 50% drop is incorporated into the model.s

Figure B3: Number of New Mortgage Contracts and Projection of Housing Loans

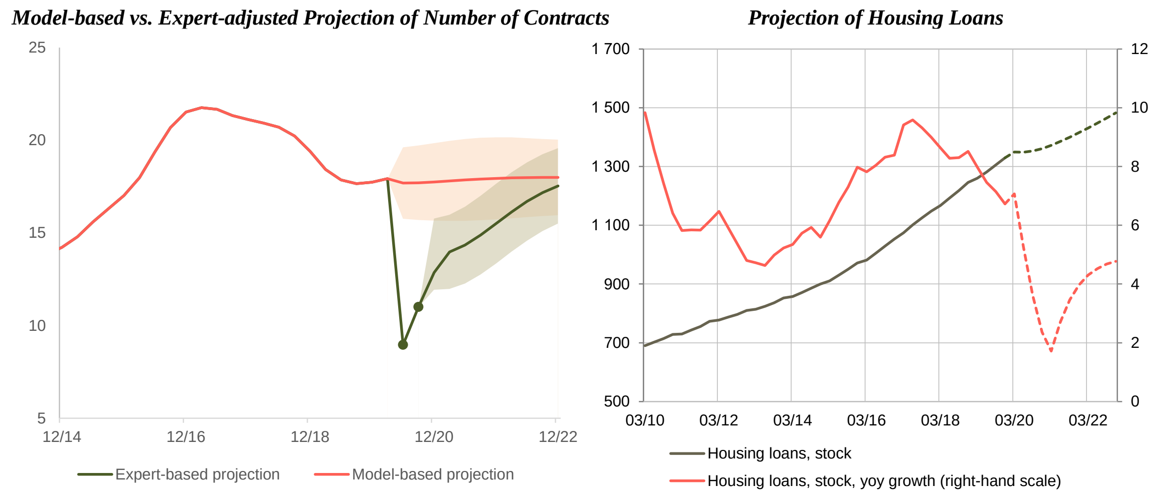

Note: Number of new contracts in thousands (left panel). The stock of housing loans in CZK billions (right panel). The dark green circles in the left figure denote points fixed on the grounds of extraneous information. The GPR regression assumes a (fictitiously) "observed" steady state of 18,000 in 2023 Q1. The shaded area (left panel) corresponds to 90% credible intervals.

The repayment schedule is based on equations (8)–(9). Since very little was known about the statutory loan moratorium at the time the scenario was designed, its first version did not include any changes to regular instalments. Later, the schedule was updated with statistical data from the newly implemented reporting by making an assumption of a roughly constant share of mortgage loans under the moratorium until its end in October 2020. In practice, the model-based repayments were proportionally adjusted for the share of loans in moratorium.[22] Optionally, the repayments could also have been adjusted for the expected rise in newly defaulting loans. A default rate estimate is available from the stress test of the mortgage loan portfolio (Hejlová and Gregor, 2020). However,

[21] In retrospect, these estimates turned out to be very pessimistic and did not even get close to reality. However, at that time they best reflected the available information and the economic narrative.

[22] Identical results could have been achieved by considering the effect of the moratorium on the average maturity of the outstanding mortgage stock.

it was assessed to rise only marginally despite the worsening economic conditions. Its omission should thus be barely noticeable in the overall prediction error.

The new credit volumes in tandem with the repayment of outstanding loans make it possible to generate projections for the stock of housing loans (Figure B3, right panel). The projection implies a considerable slowdown in year-on-year dynamics throughout 2020, with a pick-up in growth in the second half of 2021. The growth remains in positive territory over the whole forecast horizon, showing that the long-term nature of mortgage contracts makes the series remarkably persistent even under a massive decline in the provision of new credit. Again, this is something that the older generation of reduced-form satellite models completely missed, as they exhibited highly unrealistic growth rates of below -5% even under much milder economic conditions. Although the actual growth in housing loans considerably outpaced the projected growth, we recall again that the projection should above all remain consistent with the message and quantitative features of the official CNB forecast (which deviated from the subsequently observed economic developments).

B.3 Loans to Non-financial Corporations

The CNB's official spring 2020 forecast already accounted for a considerable slowdown in investment activity during 2020. This immediately translated into the projection of the loan stock through the credit impulse (Section 3.1). However, several considerations needed to be made prior to delivering the final projection, because a number of factors were pointing to a higher outstanding stock than the credit impulse alone suggested.

First, given the share of foreign currency loans, the predicted exchange rate depreciation of 7% would have increased the stock denominated in Czech koruna by approximately 2–3%. Second, the expected rise in default rates (9.5% cumulatively over the three-year horizon) would have contributed to slower debt repayment, as would the statutory loan moratorium, which allowed for temporary postponement of loan instalments (although little was known about the moratorium and its use by non-financial corporations at the time). Slower debt repayment would effectively have reduced the downward pressure on the loan stock implied by the credit impulse. Finally, it was expected that many corporations would resort to increased short-term financing in order to cover their operating expenses during the lockdown of the economy. The need for operational financing was expected to remain elevated even in the recovery phase, as many financially exhausted companies would have faced substantial financial stress otherwise. A rough estimate of the operational financing needs was obtained using a macro module of the stress-testing simulation tool (Siuda, 2020). Using the above information, the outcome of the GPR model with the credit impulse was judgmentally adjusted upwards. This eventually resulted in a 6.8% difference at the far end of the projection compared with the model-based outcome (Figure B4).

Figure B4: Projection of Bank Loans to Non-financial Corporations

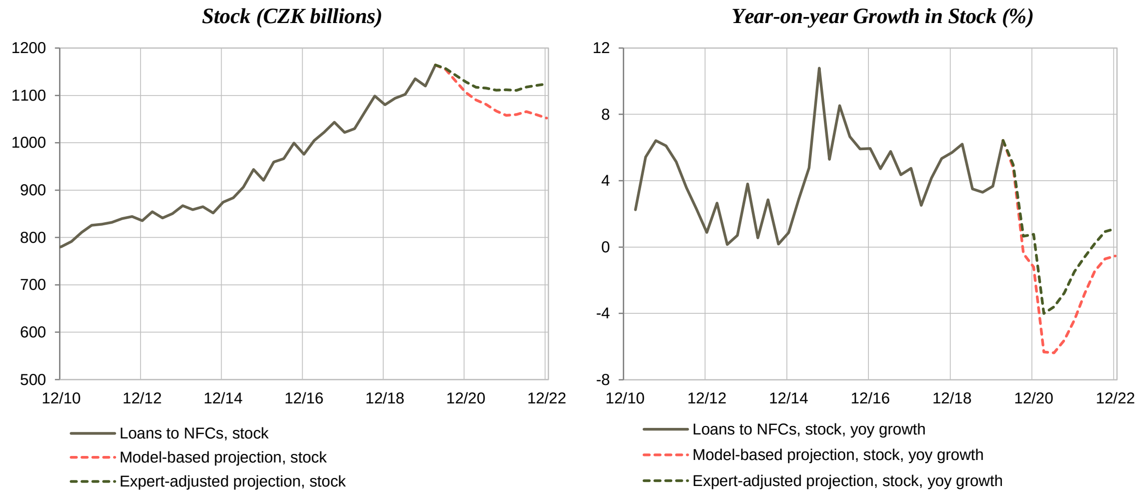

CNB Working Paper Series (since 2020)

WP 4/2021	Tomáš Adam Ondřej Michálek Aleš Michl Eva Slezáková	*The Rushin index: A weekly indicator of Czech economic activity*
WP 3/2021	Michal Franta Jan Libich	*Holding the economy by the tail: Analysis of short- and long-run macroeconomic risks*
WP 2/2021	Jakub Grossmann	*The effects of minimum wage increases in the Czech Republic*
WP 1/2021	Martin Časta	*Deriving equity risk premium using dividend futures*
WP 15/2020	Petr Polák Nikol Poláková Anna Tlustá	*How bad are trade wars? Evidence from tariffs*
WP 14/2020	Dominika Ehrenbergerová Josef Bajzík	*The effect of monetary policy on house prices – How strong is the transmission?*
WP 13/2020	Zuzana Rakovská	*Composite survey sentiment as a predictor of future market returns: Evidence for German equity indices*
WP 12/2020	Vojtěch Siuda	*A top-down stress-testing framework for the nonfinancial corporate sector*
WP 11/2020	Hana Hejlová Michal Hlaváček Blanka Vačkova	*Estimating commercial property price misalignment in the CEE countries*
WP 10/2020	Zuzana Rakovská Dominika Ehrenbergerová Martin Hodula	*The power of sentiment: Irrational beliefs of households and consumer loan dynamics*
WP 9/2020	Ivan Sutóris	*The intertemporal cost of living and dynamic inflation: The case of the Czech Republic*
WP 8/2020	Martin Hodula Jan Janků Martin Časta Adam Kučera	*On the determinants of life and non-life insurance premiums*
WP 7/2020	František Brázdik Tibor Hlédik Zuzana Humplová Iva Martonosi Karel Musil Jakub Ryšánek Tomáš Šestořád Jaromír Tonner Stanislav Tvrz Jan Žáček	*The g3+ model: An upgrade of the Czech National Bank's core forecasting framework*
WP 6/2020	Volha Audzei Jan Brůha	*A model of the Euro Area, China and the United States: Trade links and trade wars*

WP 5/2020	Dominika Ehrenbergerová Martin Hodula Zuzana Rakovská	*Does capital-based regulation affect bank pricing policy?*
WP 4/2020	Alexis Derviz	*Sovereign capital, external balance, and the investment-based Balassa-Samuelson effect in a global dynamic equilibrium*
WP 3/2020	Milan Szabo	*Growth-at-risk: Bayesian approach*
WP 2/2020	Martin Hodula Ngoc Anh Ngo	*Finance, growth and (macro)prudential policy: European evidence*
WP 1/2020	Michal Franta Ivan Sutóris	*Dynamics of Czech inflation: The role of the trend and the cycle*

CNB Research and Policy Notes (since 2020)

RPN 1/2021	Miroslav Plašil	*Designing macro-financial scenarios: The New CNB framework and satellite models for property prices and credit*
RPN 3/2020	Simona Malovaná Martin Hodula Zuzana Rakovská	*Researching the research: A central banking edition*
RPN 2/2020	Simona Malovaná Josef Bajzík Dominika Ehrenbergerová Jan Janků	*A prolonged period of low interest rates: Unintended consequences*
RPN 1/2020	Simona Malovaná	*How to organize research in central banks: The Czech National Bank's experience*

CZECH NATIONAL BANK
Na Příkopě 28
115 03 Praha 1
Czech Republic

ECONOMIC RESEARCH DIVISION
Tel.: +420 224 412 321
Fax: +420 224 412 329
http://www.cnb.cz
e-mail: research@cnb.cz

ISSN 1803-7097