Low probability, high impact: policy making and extreme events

Matthieu Bussière* & Marcel Fratzscher*

September 2004

Abstract

The objective of this note is to derive some general results for the optimal design of Early Warning System (EWS) models to anticipate extreme events, such as currency crises in emerging markets. We show how the design of an "optimal" model for policy makers focuses on the choice of three parameters: the degree of risk aversion, the forecast horizon of the model, and the probability threshold for extracting warning signals. Based on a representative EWS model, we find that, for a given degree of risk aversion, there is a unique combination of the forecast horizon and of the probability threshold that maximizes the policy maker's preferences, yielding the best possible model from a policy perspective.

JEL classification: F47, F30, D61, D78, D81, P16.
Keywords: extreme events; currency crises; early warning system; optimal design; crisis prediction.

---

* matthieu.bussiere@ecb.int and marcel.fratzscher@ecb.int, European Central Bank, Kaiserstrasse 29, D – 60311 Frankfurt, Germany. The views expressed in this paper are those of the authors and do not necessarily reflect those of the European Central Bank.

1 Introduction

The financial crises that regularly hit emerging market economies share with other extreme events – such as stock market crashes, natural catastrophes, or terrorist attacks – a number of characteristics that render them particularly challenging for policy makers. First, the social and economic costs associated with such crises are often devastating: unlike exchange rate fluctuations observed in more normal times, which usually have a limited effect on the economy, the magnitude of the collapses recorded during, for instance, the 1994 Tequila crisis or the 1997 Asian crisis have had highly disruptive effects on the domestic economies of these regions. Second, financial crises are rare events. Although such a low unconditional probability is of course a fortunate feature, it has the drawback of making crises appear unlikely to the general public, which in turn could make the decision to take preventive measures difficult politically. As policy makers usually need to take action well before the first manifestations of a crisis, they are confronted them an unpleasant trade-off: if they do not take action and a crisis happens, they can be accused of lacking foresight, but if they do and do so successfully, they may be blamed for taking apparently unnecessary measures. A third key characteristic of extreme events is that the preventive measures that need to be implemented – such as reforming the financial system or borrowing international reserves – are often themselves very costly.

The purpose of the note is not to add yet another Early Warning System (EWS) model to the already extensive literature,[1] but instead to provide some insight from a policy perspective into how to design an “optimal” EWS model that addresses the above-mentioned characteristics. The common feature of EWS models in the literature is that they consider a battery of early warning indicators (macroeconomic and financial variables) and aggregate this information in order to extract the probability of a crisis. The higher the probability of a crisis, the stronger is the case for action. But there is a need to define a cut-off point, i.e. a probability threshold above which the potential cost of a crisis outweighs the cost of taking policy action, as well as to determine a time horizon over which an extreme event is supposed to be anticipated. All other EWS models, to our knowledge, chose these parameters arbitrarily. The intended contribution of this note is to show that the choices of these parameters fundamentally alter the characteristics of EWS models.

More specifically, we take an already existing model of currency crises (Bussière and Fratzscher, 2004) and determine the optimal design of this model for policy purposes. More precisely, this involves choosing three parameters: the degree of risk aversion of missing a crisis, the probability threshold for extracting crisis signals, and the forecast horizon of the model. This last issue arises because it is more feasible to predict the time horizon rather than the exact timing of when a crisis may happen. We show that, once the degree of risk aversion is chosen, there is a unique solution to selecting the threshold and the time horizon that maximize the policy maker’s objective function.

2 The EWS model and the tradeoff problems

The standard way in the literature for constructing an EWS model is to try to explain the occurrence of a crisis (Y=1), as opposed to no crisis happening (Y=0), with a vector of explanatory variables X_j. The discrete-dependent variable model using a logistic distribution defines the logit model:

$$\Pr(Y = 1) = F(X_j \beta) = \frac{e^{X_j \beta}}{1 + e^{X_j \beta}} \tag{1}$$

A crucial element of EWS models is the attempt to anticipate the occurrence of a crisis *over a particular time horizon H* in the future. The dependent variable Y_t is therefore defined as

$$Y_t = \begin{cases} 1 & if \quad \exists k = 1 ... H \qquad s.t.\ CC_{t+k} = 1 \\ 0 & \quad otherwise \end{cases} \tag{2}$$

with *k* being the number of periods ahead of which one would like to anticipate a crisis ($CC_{t+k} = 1$). The logit model yields a predicted probability with which a crisis occurs over the chosen time horizon H. The next step consists in defining a threshold T above which the predicted probability is interpreted as a signal of an impending crisis, i.e. which implies transforming the continuous probability into a discrete variable. One key issue is how to choose the "optimal" threshold level T*. The lower it is chosen, the more signals the model will send, but at the cost of issuing more wrong signals (Type 2 errors). On the other hand, raising the threshold lowers the number of false alarms, but at the expense of missing more crises (Type 1 errors).

In essence, the preference between Type 1 and Type 2 errors determines not only the choice of T*, but also of the choice of the time horizon, or lead time, H. The following section analyses these trade-offs, using the EWS model developed by Bussière and Fratzscher (2004), which is based on a multinomial logit model proposed by Hausman and McFadden (1984).

3 The "optimal" design of EWS models

The purpose of this section is to show how changes in the preferences of the policy-maker can fundamentally alter the design of the optimal EWS model. For the purpose of this analysis, the EWS model by Bussière and Fratzscher (2004) is re-estimated over all combinations of T and H.

---

[1] Some of the seminal papers of the literature are Berg and Pattillo (1999), Eichengreen et al. (1995), Frankel and Rose (1996), Kamin (1999), Kaminsky et al. (1998), Kaminsky and Reinhart (1999). Berg et al. (2004) provide a survey and assessment of the literature.

Figure 1: Comparing differences in risk-aversion in loss functions

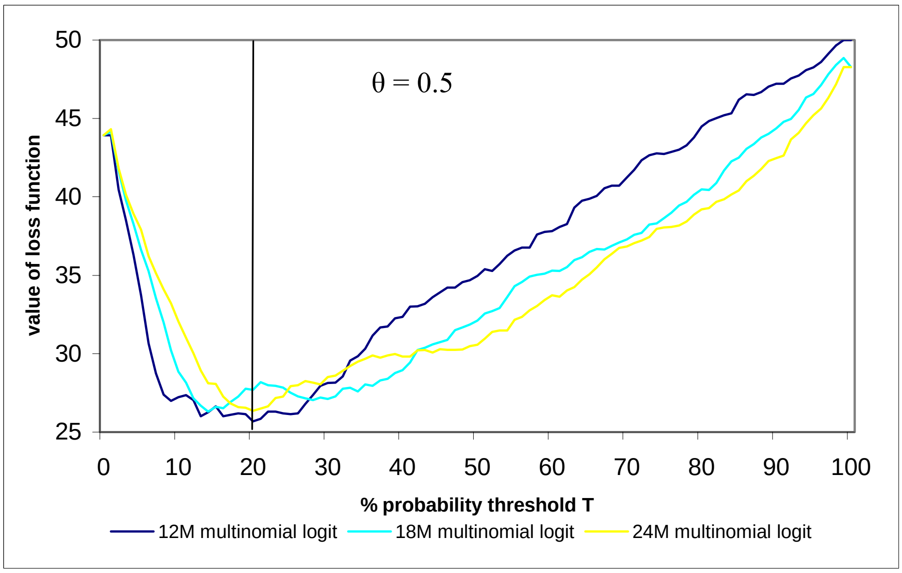

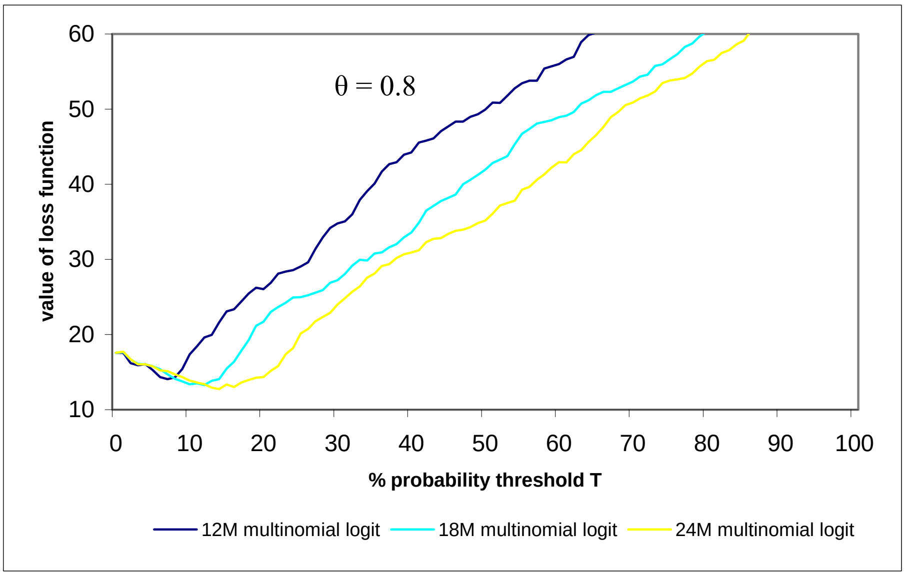

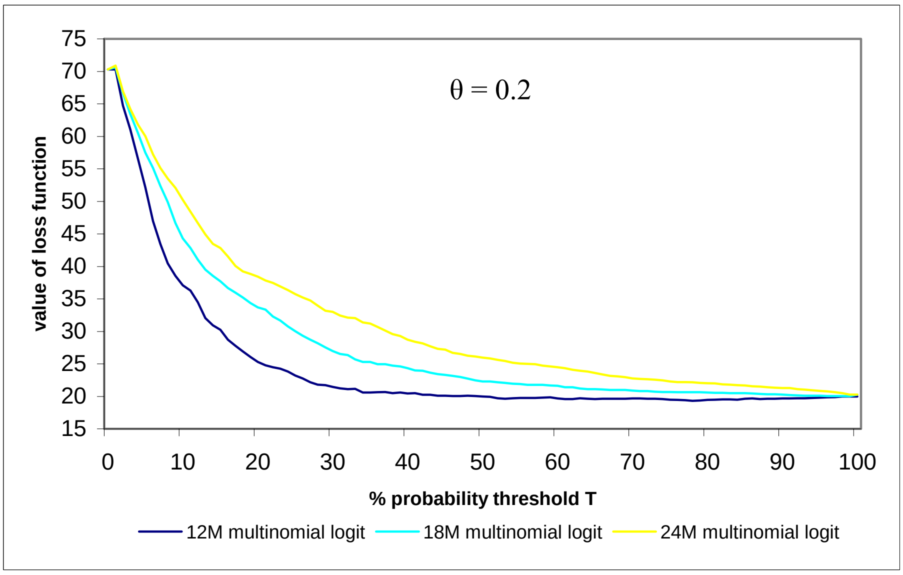

As discussed in section 2, a policy-maker wishes to simultaneously minimize the number of false alarms and the number of missed crises (Type 1 and Type 2 errors), while getting these signals as early as possible. The difficulty with the latter is that there is a trade-off between the number of false alarms and the number of missed crises: raising the probability threshold T reduces the number of false alarms, but at the cost of raising the number of missed crises, and vice versa. Changing the time horizon H leads to a similar trade-off.

Which of the models and thresholds a policy-maker wishes to choose depends on her objective function. If missing a crisis and taking pre-emptive action *both* have a cost for the policy-maker, her loss function may be formulated as

$$L(T) \equiv \theta \ \left(prob^{NS/C}(T)\right) + (1-\theta) \ \left(prob^{S}(T)\right) \tag{3}$$

with $prob^{NS/C}$ as the probability of a missed crisis, i.e. the joint probability that the EWS model issues no signal and a crisis occurs, and $prob^{S}$ as the probability of issuing a signal that a crisis will occur. θ can be interpreted as the relative cost of missing a crisis, or as the policy-maker's degree of relative risk-aversion, whereas $(1 - \theta)$ can analogously be understood as the cost of taking pre-emptive action.

Given this loss function, what threshold T* and time horizon H* should the policy-maker choose? Figure 1 illustrates how the choice of T and H is determined by θ, i.e. by the relative importance the policy-maker attaches to missing a crisis. For $\theta = 0.5$, the model that produces the lowest value of the loss function L is the EWS model with horizon $H=12$ and a threshold of $T=20$. This implies that at $T=20$, and given $\theta = 0.5$, the 12-month multinomial logit model provides the best achievable trade-off between missing crises and issuing wrong signals (Figure 1, first panel). On the one hand, lowering T raises the number of false alarms disproportionally more than the number of correctly anticipated crises. On the other hand, raising T reduces the number of false alarms, but this is more than outweighed by raising the number of missed crises. By contrast, a highly risk-averse policy-maker with $\theta = 0.8$ prefers the 24-month EWS model with a threshold of $T=14$ (Figure 1, second panel), whereas a policy-maker with $\theta = 0.2$ chooses the 12-month EWS model with a very high threshold, $T = 78$ (Figure 1, third panel).

The central result is that for any given value of θ, the loss function admits an absolute minimum which determines both the optimal time horizon H* and the optimal threshold T*. Overall, estimating our multinomial model for all possible degrees of risk aversion θ, for all possible thresholds T and for all time horizons H between one month and 24 months, allows us to derive three general results:

Result 1: For any given degree of risk aversion θ, the larger the time horizon H* of the optimal model, the larger is the optimal threshold T* (see Figure 2).[2] Indeed, at any given threshold T, models with a longer time horizon H yield a relatively higher proportion of false alarms than correct signals. This is because the predictive power of the EWS model becomes lower, for any given T, the longer the lead time with which a crisis is supposed to be anticipated. Therefore, a longer time horizon H requires also the raising of the threshold T in order to minimize the loss function at any given θ.

Result 2: For any given time horizon H, raising the degree of risk aversion θ leads to a lower optimal threshold level T* (see Figure 2). The reason is that the higher the weight a policy-maker puts on not missing a crisis, the lower she will wish to set the threshold T of the model in order to raise the number of correct crisis signals.

Result 3: For any given probability threshold T, a higher degree of risk aversion θ induces a longer optimal time horizon H* (see Figure 3). The intuition is that a longer time horizon *H* allows a policy-maker to miss fewer crises, which outweighs the fact that choosing a longer time horizon *H* also raises the number of alarms that the EWS model sends. In the example with three time horizons (*H = 12, H = 18, H = 24*), the multinomial logit EWS model with *H = 12* performs best for $0 \leq \theta < 0.55$, the EWS model with *H = 18* is best for $0.55 \leq \theta < 0.65$, and the EWS model with *H = 24* is the preferred model for $0.65 \leq \theta < 1$.

Figure 2: Determining the optimal threshold T*

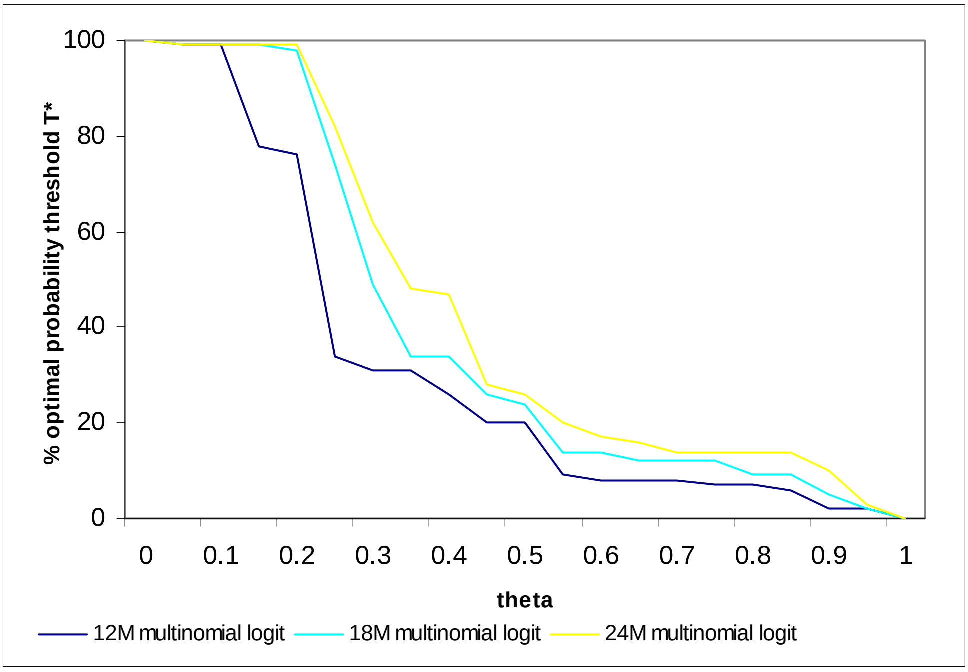

[2] T* is determined by estimating each multinomial logit model with time horizon H over all $\theta \in [0,1]$ and all thresholds $T \in [0,1]$. T* that minimizes the loss function at each θ is then graphed in Figure 2.

Figure 3: Determining the optimal time-horizon H*

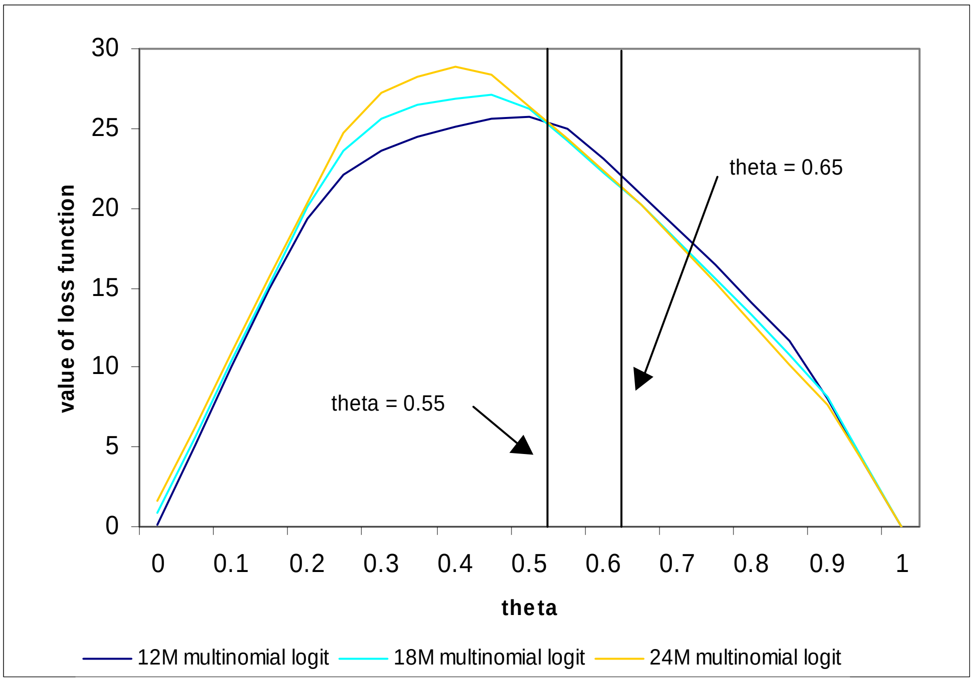

4 Conclusions

The objective of this note has been to provide some general results for the optimal design of EWS models to anticipate extreme events, such as currency crises. The central argument of the paper is that the design of an "optimal" EWS model focuses on the choice of three parameters: the degree of risk aversion θ of missing a crisis, the forecast horizon H of the model, and the probability threshold T for extracting crisis signals. The empirical results indicate that, first, a higher degree of risk aversion should induce modelers to choose a longer time horizon H and a lower threshold T. Second, for any given degree of risk aversion, a choice of a longer time horizon H should optimally require a higher threshold T, and vice versa. Among these three parameters, the degree of risk aversion can be considered as exogenous for the modeler, since it depends on (i) the relative cost of a crisis, compared with the cost of taking preemptive action and (ii) the policy maker's preferences. Once this parameter has been established, a central result of the note is that the time horizon H and the threshold T are uniquely determined, yielding a single "optimal" EWS model from a policy perspective. Overall, these implications can easily be extended to other circumstances where policy makers have to deal with extreme events, and help modelers to construct economic models that are more realistic and closer to the objectives of policy makers.

References

Berg, A., Borensztein E., Pattillo, C., 2004. Assessing early warning systems: how have they worked in practice? IMF Working Paper 04/52.

Berg, A., Pattillo, C., 1999. Are currency crises predictable? A Test. IMF Staff Papers 46(2), 107-38.

Bussiere, M., Fratzscher, M., 2004. Towards a new early warning system of financial crises. Journal of International Money and Finance (forthcoming) and ECB Working Paper No. 145, May 2002.

Eichengreen, B., Rose, A., Wyplosz, C., 1995. Exchange market mayhem: the antecedents and aftermaths of speculative attacks. Economic Policy 21, 249-312.

Frankel, J., Rose, A., 1996. Currency crashes in emerging markets: an empirical treatment. Journal of International Economics 41, 351-66.

Hausman, J., McFadden, D., 1984. A specification test for the multinomial logit model. Econometrica 52, 1219-1240.

Kamin, S., 1999, The current international financial crisis: how much is new? Journal of International Money and Finance 18(4), 501-514.

Kaminsky G., Lizondo, S., Reinhart, C., 1998. Leading indicators of currency crises. IMF Staff Papers 45(1), 1-48.

Kaminsky G., Reinhart, C., 1999. The twin crises: the causes of banking and balance-of-payments problems. American Economic Review 89(3), 473-500.