Review of Economic Dynamics 18 (2015) 212–226

Contents lists available at ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red

Optimal stopping in a model of speculative attacks

Pablo Kurlat [*,1]

Stanford University, United States

ARTICLE INFO

Article history:
Received 24 March 2013
Received in revised form 14 March 2014
Available online 1 April 2014

JEL classification:
D82
D84
E58
F31

Keywords:
Speculative attacks
Global games
Currency crises
Bank runs

ABSTRACT

When faced with a speculative attack, banks and governments often hesitate, attempting to withstand the attack but giving up after some time, suggesting they have some ex-ante uncertainty about the attack they will face. I model that uncertainty as arising from incomplete information about speculators' payoffs and find conditions such that unsuccessful partial defences are possible equilibrium outcomes. There exist priors over the distribution of speculators' payoffs that can justify any possible partial defence strategy. With normal uncertainty, partial resistance is more likely when there is more aggregate uncertainty regarding agents' payoffs and less heterogeneity among them.

1. Introduction

In August and September of 1992, the Bank of England sold billions of dollars of foreign reserves in an attempt to maintain the pound's exchange rate within the bands of the European ERM. On September 16 it finally gave up and abandoned the ERM. The cost of having attempted to defend the parity was later estimated at approximately £ 3.3 billion (UK Treasury, 1997).

This pattern, of first attempting to defend the existing regime and giving up after some time, is a common feature of speculative attacks and explaining it presents a theoretical challenge. So-called first-generation models based on Krugman (1979), such as Flood and Garber (1984) and Broner (2007) account for it in a very simple way: by assuming the government follows and attempts to defend an unsustainable policy for exogenous reasons and abandons it only when forced to do so. However, these models leave unanswered the question of why a government would behave this way.

As formalized by Obstfeld (1996) and others, speculative attacks often have a self-fulfilling aspect: if enough agents believe the government will abandon a regime, they will act in ways that make it optimal for the government to indeed abandon it. The unsatisfying conclusion of models of self-fulfilling equilibria is that, at least within some range of parameters, the outcome is arbitrary or depends on ad-hoc unmodeled factors. Following Morris and Shin (1998), many authors have argued that modifying the common-knowledge assumptions of games that have self-fulfilling equilibria may help to resolve this indeterminacy and provide more definite predictions.

* Correspondence to: Department of Economics, Stanford University, Stanford, CA 94305, United States.
E-mail address: pkurlat@stanford.edu.

1 I am deeply indebted to George-Marios Angeletos, Ricardo Caballero, Bengt Holmström and especially Iván Werning for their invaluable guidance. I also thank Daron Acemoglu, Sergi Basco, Emre Ozdenoren and seminar participants at the MIT Macro Lunch for useful suggestions and comments. All remaining errors are my own.

http://dx.doi.org/10.1016/j.red.2014.03.003

The basic building block of models in this literature is a game played by many small agents ("speculators") who are incompletely informed about the relevant parameters of the economy, and one large agent ("the central bank") who has complete information. Typically, the focus of the analysis is on the structure of actions and information of the game played by the speculators. In contrast, the central bank's information and objectives are usually described in very simple terms so that its strategy can be summarized, or even replaced, by a simple rule such as "defend the existing regime unless a mass of speculators larger than A^* attacks it".

This paradigm (and for that matter the multiple-equilibria paradigm too) fails to account for why the central bank, acting rationally, would ever engage in an unsuccessful partial defence of the regime, as the Bank of England did in 1992. In these models, the central bank knows the "fundamentals" of the economy and can therefore perfectly predict (or in some versions observe) what the size of the speculative attack is going to be. As long as defending the regime is costly, it would never be the case that it attempts to defend it but surrenders after some time, since this failure would have been foreseen.

However, in some contexts the possibility of a temporary, unsuccessful defence of the status quo makes an important difference. For example, if we wish to apply these methods to the study of bank runs, as Goldstein and Pauzner (2005) do, the *only* reason why depositors would run is if they believe that the bank will pay some of them before falling or deciding to suspend convertibility.

How can the theory account for the phenomenon of unsuccessful defences? One possibility, implicitly assumed by Morris and Shin (1998), is that defending the regime is not costly at the margin; conditional on regime change, the central bank does not have a preference for how far it held out. In many contexts this is not a reasonable assumption: the losses to the central bank's balance sheet are greater the more reserves it has spent trying to defend a fixed exchange rate; the liquidation costs a bank incurs in are smaller the sooner it suspends convertibility; the retaliation against a dictator is likely to be harsher the longer he held on to power; the casualties to a surrendering garrison are greater the longer it attempted to withstand a siege.

If unsuccessful defences are costly, a theory that accounts for them must somehow allow the central bank to have uncertainty about the size of the attack it is going to face.[2] With uncertainty and suitable timing assumptions, the central bank's decision may be viewed as an optimal stopping problem: as the attack escalates, it must decide whether to surrender or to continue to defend the regime in the hope that the attack will be over soon, using its appropriately updated beliefs about how large the attack is likely to be.

There is more than one way to introduce uncertainty into the central bank's problem. One approach, pursued by Ennis and Keister (2010) is to let speculators play a correlated (sunspot) equilibrium, and assume that the central bank does not observe the sunspot. A related possibility is to abandon pure-strategy Nash equilibrium as a solution concept. In any pure-strategy Nash equilibrium, the central bank knows the strategies of the speculators, and is thus able to predict the size of the speculative attack with no uncertainty. However, under appropriate conditions (although not in the Morris–Shin limit), both attacking a regime and not attacking it are rationalizable actions. If the requirement that the central bank know the speculators' strategies is dropped, it is possible to simply endow it with beliefs about the joint distribution of (rationalizable) actions the speculators might take, and under these beliefs a policy of partial defence may indeed be optimal. The trouble with explaining the phenomenon along either of these lines is that these explanations rely on arbitrary assumptions about the central bank's beliefs and/or speculators' reactions to the sunspot.

This paper introduces uncertainty into the central bank's decision problem in a different way. As in Diamond and Dybvig (1983), Wallace (1988), Green and Lin (2003) and Peck and Shell (2003), there is aggregate uncertainty about the distribution of (heterogeneous) preferences in the population of speculators. This distribution is governed by a single random parameter θ, and neither the central bank nor the speculators know its realization. Although the central bank knows the equilibrium strategies, the equilibrium size of the attack, conditional on the central bank's information, is a random variable, so it faces a nontrivial optimal stopping problem.

For the central bank's optimal stopping problem to have an interior solution, in which the central bank surrenders after some time despite having incurred sunk costs of defence, it must be that as the attack progresses the central bank becomes sufficiently more pessimistic about the magnitude of the attack that it will face. In particular, it must think that an attack that is not over by the stopping point is unlikely to be over soon after that, which is consistent with some ex-ante beliefs $f(A)$ about the size of the attack A but not with others. However, I show that given any probability distribution $f(A)$, it is possible to reverse-engineer a prior about θ such that $f(A)$ is indeed the endogenous probability distribution in an equilibrium of the game. Hence a basic finding is that the model is able to deliver the kinds of uncertainty that could justify partial defences.

This alone, however, provides no clear guidance as to what factors make unsuccessful defences likely to arise. To answer this question, I specialize the model to a simple example with linear payoffs and normal uncertainty. In this case, I show that partial defences are more likely to arise and last longer if they do arise if heterogeneity in preferences is small relative to aggregate uncertainty about average preferences. The model has multiple equilibria for some parameter values. The central bank's uncertainty and the possibility of partial unsuccessful defences, however, are not due to this but to uncertainty about outcomes *within* a given equilibrium.

[2] This is also noted by Goldstein et al. (2008). In their model the central bank has uncertainty about the value of maintaining a fixed exchange rate and may learn about this by observing the speculative attack. Their model does not, however, allow the central bank to surrender to the attack midway.