Discussion of Alvarez and Dixit: A Real Options Perspective on the Euro

David Backus
Stern School of Business, New York University, and NBER
Email: dbackus@stern.nyu.edu
Phone: 212 998 0873
Address: NYU Stern School of Business, 44 West Fourth Street, New York NY 10012

November 21, 2013

Abstract

Discussion of a paper by Fernando Alvarez and Avinash Dixit presented at the Carnegie-Rochester-NYU Conference on Public Policy, held at New York University on April 19-20, 2013.

JEL Classification Codes: F31, F33, C61.

Keywords: the euro; Euro Area; real option; option value.

Comments on Alvarez and Dixit

As I told Mark and Hal, I'm delighted to talk about this paper. It's on an interesting topic — the possible breakup of the Euro Area — and the authors are two of my favorite economists. I'll describe what they do, then talk about how it informs our view of what's going on in Europe right now.

Perpetual options

Since the Alvarez-Dixit model has a similar mathematical structure, I thought I'd start by reviewing perpetual options: options that, unlike most financial market examples, have no expiration date. Valuation takes a beautiful recursive form, as we decide each period whether to exercise the option, which we can do only once, or wait another period. I find the logic incredibly clear in discrete time, so I'll either clarify or run roughshod over the paper's continuous-time math, depending on your point of view.

Consider asset pricing in a stationary Markov setting with a state variable x. The ex-dividend value of a claim to the stream of future dividends d might be expressed

$$V(x_t) \quad = \quad E_t\Big\{m(x_t, x_{t+1})\big[d(x_{t+1}) + V(x_{t+1})\big]\Big\}, \tag{1}$$

where m is the pricing kernel. The value of a perpetual option to buy this asset at strike price k is then

$$J(x_t) \quad = \quad \max\Big\{E_t\left[m(x_t, x_{t+1})J(x_{t+1})\right], V(x_t) - k\Big\} \tag{2}$$

The right branch of this Bellman equation is the value of exercising the option now, the difference between the market price and the strike. The left branch is the value of waiting till next period, discounted back to the present.

The solution has a number of typical features, some of which require additional structure:

- Threshold property. The solution has the form: exercise if $V(x_t) \geq V^*$ for some threshold value V^*, wait otherwise.

- Convexity. Options have convex payoffs. The one-period payoff $\max\{0, V-k\}$ is convex in V. If we rewrite the problem so that V is the state variable, this leads to a convex value function J.

- Option value. One consequence of convexity is that there's value in waiting: generally $V^* >> k$, which means we wait for V to rise well above the strike before exercising the option.

- Volatility. Another consequence of convexity is that the value of the option increases with uncertainty. A mean-preserving spread, for example, raises $J(V)$ and V^*. Why? Because there's a greater chance we'll get lucky. There's also a greater chance we'll get unlucky, but the option chops off the left tail.

I give an example in the appendix. All of these features show up, in one form or other, in the Alvarez-Dixit model.

The Alvarez-Dixit model

Their model captures some of the salient features of the common currency of the Euro Area. One feature is the benefit of a common currency. That shows up here as a constant positive payoff every period the system is in place. Another feature is the cost of imposing the same monetary policy on every country. That shows up here as squared deviations from purchasing power parity. I think we want to interpret these deviations flexibly, so I'll refer to them simply as deviations.

Here are the ingredients. Each country i has a state variable X_{it}, an AR(1) with normal innovations. With policy Z_{it}, the deviation is $x_{it} = X_{it} - Z_{it}$. The welfare of country i is

$$u_i \;=\; \begin{cases} -x_{it}^2 & \text{with independent policy} \\ \alpha - x_{it}^2 & \text{with common policy,} \end{cases}$$

where $\alpha > 0$ is the benefit of a common currency. Aggregate welfare is the sum. With common policy, that's

$$U \;=\; \sum u_i \;=\; n\alpha - \sum x_{it}^2.$$

With independent policies, each country sets $Z_{it} = X_{it}$ so that the deviation x_{it} is zero. Welfare is zero, both individually and in the aggregate. With common policy, the optimal policy sets the average deviation equal to zero with $Z_t = n^{-1} \sum_j X_{jt}$. The question is whether welfare is greater with common policy, which contributes $n\alpha$ to aggregate welfare but generates deviations that reduce welfare.

They introduce a breakup option that mirrors the perpetual option problem. If they (meaning the Euro Area as a whole) pay a breakup fee of nk, they can dissolve the common currency system and revert to individual country policies, in which welfare is zero. (They label the fee Φ, but k seems to me a better fit for an option.) Here's how that works. There's one really clever trick here, which is to express aggregate welfare in terms of a single state variable,

$$Y_t \;\; = \;\; \sum x_{it}^2.$$

The same trick is used in Fernando's earlier work on price setting with Francesco Lippi (Alvarez and Lippi, 2013). The Bellman equation for the breakup option is then

$$J(Y_t) \;\; = \;\; \max\Big\{ n\alpha - Y_t + e^{-r} E_t\big[J(Y_{t+1})\big], 0 - nk \Big\}$$

The right branch is the breakup option: pay the fee nk and revert to the welfare of zero you get from following individual country policies. The left branch is the value of staying in the common currency system for another period. Each country then gets the benefit α minus the cost of deviations, now summarized by Y_t. Future value is discounted by e^{-r}.

The solution, which they find numerically, has familiar features: the threshold property, option value, and so on. They're described in numerical examples, designed to be plausible. One difference from the traditional option problem is that the impact of volatility is ambiguous. Why? I think the answer is that increasing volatility of the X's increases the mean as well as the volatility of Y.

What does this tell us about the Euro Area?

Let's step back and think about what's going on in the doomsday machine that is Europe today. What do we learn from this model? How do we interpret it? What have we missed?

How should we think about deviations? The authors suggest that deviations are departures from purchasing power parity. Using numbers from flexible exchange rate regimes, they choose parameters that generate a standard deviation of annual exchange rate changes of about 8%, which is roughly what you'd see for the US dollar against the euro, the yen, or the pound.

Is that the right comparison? My first thought is that this is way too high. One thing the Euro Area has clearly accomplished is a sharp reduction in real exchange rate volatility. For the Euro Area, something like 1% annually would fit the evidence better.

But that's probably wrong. What I've learned over the years, much of it from Enrique Mendoza, is that crises generate large real exchange rate movements — much larger than the deviations in the paper. In Mexico in 1994, the peso fell 30% in a month. In Korea in 1997, the won fell 40%. In Argentina in 2002, the peso fell 65%. If we interpret these events as telling us how much adjustment is needed in a crisis, we'd be justified in using much larger values for the volatility of deviations. (We could also add jumps, which I think is realistic, but they would call for more sophisticated methods than the already sophisticated methods in the paper.) It's not clear what this would do to the examples, given the ambiguity of the impact of volatility, but it's worth thinking about.

Collective v. individual action. Another issue is the distinction between individual countries and the Euro Area. The paper starts by describing how the Euro Area would maximize aggregate welfare given a breakup option. Since breakup is irreversible, there's value to waiting. In numerical examples, the common currency is expected to last for decades.

However, political power in Europe lies predominantly in countries. One way to think about the implementation of the collective decision of the Euro Area is that they make transfer payments that spread the benefits around. But suppose transfer payments are ruled out. Certainly it seems as if transfer payments are more difficult in Europe than they are in the US, so that deserves some thought. The authors consider, in this setting, the problem of a country that has a unilateral option to exit at price k. They find that exit by such a country would occur more quickly than a collective decision to break up the currency union. In that sense, transfers are critical to the continuation of the system.

What about debt? Debt has gotten most of the attention in the press, and it's missing from the model. But could we think about deviations as somehow reflecting sovereign debt problems? This would also, of course, raise new issues about transfer payments, and whether they make sense in this context. But it also suggests that we're missing a state variable that makes deviations more or less painful. Portugal and Spain, for example, has a greater interest in inflating away their debt than Germany does. So debt, I think, is something that requires us to go well beyond what we see in the model.

Where does that leave us? We have an interesting and provocative paper that elucidates some of what we see in Europe today. Fortunately for the rest of us, they've left some issues for future work.

Appendix: A perpetual option example

As I was preparing my discussion, I ran across an example of a perpetual option that can be solved by hand. My version is adapted from Gerber and Shiu (1994). It's based on the same lognormal structure that delivers the Black-Scholes-Merton formula for prices of fixed-maturity options.

Suppose the state is the price V of the underlying asset. We apply risk-neutral pricing, which you might think of as using a constant pricing kernel $m = e^{-r}$. The future price of the underlying is lognormal: $\log V_{t+1} - \log V_t \sim \mathcal{N}(\kappa_1, \kappa_2)$. The dividend is proportional to V, so that $d_{t+1} + V_{t+1} = e^{\delta} V_{t+1}$ for some parameter $\delta > 0$. With these inputs, the valuation equation (1) becomes

$$V_t \;=\; E_t\big[m_{t+1}(d_{t+1} + V_{t+1})\big] \;=\; E_t\left(e^{-r} e^{\delta} V_{t+1}\right) \;=\; e^{\delta - r}\; V_t\; e^{\kappa_1 + \kappa_2/2},$$

which implies

$$0 \;=\; \delta - r + \kappa_1 + \kappa_2/2.$$

This *no-arbitrage condition* is a consistency check on our assumptions. We'll use it to nail down κ_1.

We find the price of a perpetual (call) option by guess and verify. The option valuation equation (2) is now

$$J(V_t) \;=\; \max\left\{E_t\big[e^{-r} J(V_{t+1})\big], V_t - k\right\}.$$

We guess $J(V) = AV^a$ for parameters (A, a) to be determined. That gives us three unknowns: A, a, and the threshold V^*. If $V = V^*$ we're indifferent between exercising and not, so we have

$$J(V^*) \;=\; A\,V^{*a} \;=\; e^{-r} A\,V^{*a}\,e^{a\kappa_1 + a^2\kappa_2/2} \;=\; V^* - k.$$

That's two equations. Alvarez and Dixit refer to them as *value matching*. The third is the envelope condition,

$$J_V(V^*) \;=\; aAV^{*a-1} \;=\; 1.$$

In continuous time they call this the *high contact* or *smooth pasting condition*. The first equation gives us a:

$$0 \;=\; a^2\kappa_2/2 + a\kappa_1 - r \quad \Rightarrow \quad a \;=\; \frac{-\kappa_1 + (\kappa_1^2 + 2\kappa_2 r)^{1/2}}{\kappa_2}.$$

We take the positive root; otherwise, the value function is decreasing in the price of the underlying. The other two equations give us $V^* = [a/(a-1)]k > k$.

This example has all the properties we mentioned earlier. First, the solution has the threshold property. Second, since $a > 1$, the value function is convex. Third, option value can be large. Consider the numerical example: $r = 0.02$, $\delta = 0.01$, and $\kappa_2 = 0.1^2$. Then $\kappa_1 = 0.005$. If the strike is $k = 7$, then $a = 1.56$ and $V^* = 19.47 >> k$. Fourth, if we increase uncertainty κ_2, then a falls and V^* rises.

References

Alvarez, Fernando, and Francesco Lippi, 2013, "Price setting with menu cost for multi-product firms," *Econometrica*, forthcoming.

Gerber, Hans, and Elias S.W. Shiu, 1994, "Martingale approach to pricing perpetual American options," *Astin Bulletin* 34, 195-220.

Acknowledgements

I thank Mike Chernov and Stan Zin for conversations on related topics, including the logic behind perpetual options. I also thank, with only a touch of ambivalence, the organizers, Mark Aguiar and Hal Cole, for giving me this opportunity.