Online Appendix: Optimal Financial Exclusion

By Cyril Monnet and Erwan Quintin*

This appendix to "Optimal Financial Exclusion" discusses the consequences on our results of maximizing date-0 welfare rather than stationary welfare and the effects of heterogenous default costs on optimal exclusion length.

I. Maximizing date zero welfare

We have made the case that exclusion policies that maximize stationary welfare must front load-punishment. Does front-loading punishment also maximize the welfare of the investors who happen to be alive at date zero, with types arbitrarily distributed? If possible, it is clearly efficient to forgive all the investors in bad standing at date 0 since punishing them has no remaining impact on incentives while forgiving them makes the volume of transactions as high as it can be at date 0. We will consider both the case where those amnesties are possible and the case where they are not.

For simplicity, we will concentrate our attention once again on the situation where problem $(P1)$ has no solution so that investors in bad standing do not operate their project. Furthermore, we also continue to restrict our attention on symmetric SSEs, defined as before, with the implied constant continuation utilities. Such an SSE continues to exist for all possible forgiveness policies and the arguments for why this is so are unchanged.

In keeping with our earlier notation, denote the mass of investors in good standing at date 0 by μ_0^G. When amnesty is an option, it is efficient to set $\mu_0^G = 1$ via immediate forgiveness but, in general, the mass $\mu_0^B(n)$ of incumbent agents who, at date 0, have been excluded for n period may be positive. Efficient forgiveness policies must maximize

$$\mu_0^G V^G + \sum_{s=0}^{+\infty} \mu_0^B(s) V^B(s)$$

where V^N is the expected lifetime utility of agents who are not excluded at date 0 under the assumed SSE while $V^B(s)$ is the same for agents who have been excluded for s periods subject to

$$V^B(0) = \phi_0 V^G + (1-\phi_0)\phi_1 \beta V^G + (1-\phi_0)(1-\phi_1)\phi_2 \beta^2 V^G + \dots$$

Importantly, we no longer need to impose $\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty$. As we will

* Cyril Monnet: BIS, University of Bern and SZ Gerzensee, cyril.monnet@vwi.unibe.ch. Erwan Quintin: University of Wisconsin – Madison, equintin@bus.wisc.edu.

see below in fact, policies that imply that eventually all agents are excluded may well be optimal when maximizing date zero welfare.

To proceed, assume an amnesty of all agents who are in bad standing at date 0 is feasible. Following amnesty, the economy begins date 0 with all agents in good standing. In that case, SSEs are trivial to rank in terms of welfare. The lifetime utility of non-excluded agents at date 0 is fully summarized by the value $V^B(0)$ of becoming excluded since it implies the terms on loans and, in turn, V^G which is all we need to know when $\mu_0^G = 1$. It follows that exclusion policies that imply the same $V^B(0)$ all result in the same welfare level. This implies that multiple policy shape may now be optimal. Indeed, forgiveness policies with very different profiles may imply the same level $V^B(0)$ of initial punishment. To show that this is in fact a possibility, consider a parametric version of our environment in which agents discount the future at a rate $\beta = 0.85$, the project requires one unit of capital, the probability of success is $\pi = 0.95$, output when positive is $y = 1.5$, and the opportunity cost of funds is $R = 1.2$. The default cost τ follows a log-normal distribution with location parameter 2 and dispersion parameter 1.

Consider then a flexible, sigmoid class of forgiveness policies characterized by two parameters (a, b) such that for $n \geq 0$:

$$\phi_n = \frac{1}{1 + \exp\left[a(n-b)\right]}$$

This family can closely approximate most monotonic exclusion policies including step-functions in which case b pins down the location of the inflection point while a pins down the steepness of the inflection. The sign of a determines whether forgiveness odds rise or fall with time in exclusion. This specification thus allows both for policies that front-load and policies that back-load punishment.

In this parametric case, the two policies displayed in figure 1 imply the same equilibrium level of $V^B(0)$ hence the same level of welfare. We located these two welfare equivalent policies by searching numerically for welfare maximizing policies given $V^B(0)$ starting from different initial conditions. The search stops the moment the procedure has found a candidate that achieves the target level of $V^B(0)$. A numerical procedure is needed to look for those welfare-equivalent policies since any change in the policy implies a change in lending terms and value function. The first step of our procedure solves for equilibrium given an exclusion policy. The second step looks for policies that maximize V^G given $V^B(0)$. Our code is available for download at http://erwan.marginalq.com/index_files/wp.htm.

Many other policies, including non-monotonic ones, deliver the same welfare level but those two specific examples suffice to convey two key ideas. First, when amnesty is an option, the efficient shape of exclusion policies is indeterminate. One of the two policies displayed in the figure front-loads punishment in the sense that forgiveness odds are initially low but then rise, while the other back-loads

FIGURE 1. WELFARE-EQUIVALENT EXCLUSION POLICIES WHEN AMNESTY IS AN OPTION

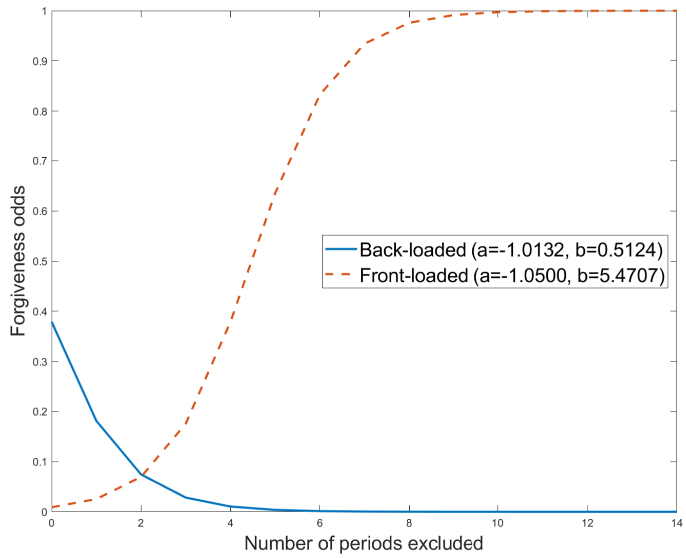

punishment. In other words, while front-loading is uniquely efficient under the long-term perspective we adopted in previous sections, it is only weakly optimal when maximizing the welfare of investors alive at date 0 and when amnesty is an option. Second, policies that violate ergodicity criterion

$$\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty$$

can be efficient. The back-loaded policy shown in the figure has forgiveness policies converge to zero. Therefore, there is a positive probability that excluded agents may never return to the non-excluded fold. The long-term distribution of types, in that case, is degenerate at $\mu^G = 0$.

Front-loading punishment is thus but one of countless efficient shapes for exclusion policies when amnesty is possible. When amnesty is not feasible, on the other hand, our efficiency criterion takes into account the welfare of existing agents who happen to be excluded at date 0. This restores benefits for front-load punishment, as the next result shows.

PROPOSITION 1: *Let n^* be the highest number of periods for which borrowers alive at date 0 have been excluded, i.e. $n^* = \sup\left\{n : \mu_0^B(n) > 0\right\}$. Then, at any efficient policy, either $\phi_{n^*} = 1$ or $\phi_n = 0$ for all $n < n^*$.*

PROOF:

Among agents alive at date 0, the welfare of investors in good standing and investors who just lost their good standing ($n = 0$) is pinned down by $V^B(0)$. Therefore, these agents are indifferent across all exclusion policies that deliver $V^B(0)$. Assume that $\mu^B(1) > 0$ and consider any policy such that that $\phi_1 < 1$ while $\phi_0 > 0$. Recall that we must have

$$\mu^B(0) = \phi_0 V^G + (1 - \phi_0)V^B(1).$$

Holding all value functions the same then, we can lower ϕ_0 while maintaining $V^B(0)$ by raising $V^B(1)$, leaving the welfare of investors in good standing and those who just lost their standing ($n = 0$) unaffected, but making agents who have been in bad standing for one period ($n = 1$) strictly better off. In addition, since $V^B(0)$ is unchanged, any set of SSE loan terms remain part of a SSE, so that V^G remains part of a SSE as well. This means that any policy such that $\phi_1 < 1$ while $\phi_0 > 0$ is suboptimal if $\mu^B(1) > 0$. This proves the result when $n^* = 2$. Extending the argument recursively establishes the proposition for all n^*. ■

The argument above begins with the same observation that leads to indeterminacy in the case with amnesty: under any SSE, the lifetime utility of agents who are in good standing at date 0 is fully summarized by $V^B(0)$. The same holds, obviously, for agents who just lost their standing ($n = 0$). Those agents are indifferent across all policies that deliver $V^B(0)$. Policies that front-load punishment ($\phi = 0$) on the other hand, benefit agents who have been excluded for more period ($n > 0$) since more punishment early means less necessary punishment later. This result has the following key implication for our purposes:

COROLLARY 2: *If $\mu_0^B(n) > 0$ for all $n \leq n^*$, then all efficient exclusion policies must have the shape described in Theorem 5 up to n^*.*

PROOF:

If $\phi_{n^*} < 1$ at all optimal policies then the result follows. If $\phi_{n^*} = 1$ then repeating the argument above implies that either $\phi_{n^*-1} = 1$ or $\phi_n = 0$ for all $n < n^* - 1$, and the result follows. ■

The premise that $\mu_0^B(n) > 0$ for all $n \leq n^*$ should generally be expected to hold since any policy that induces a stationary distribution such that $\mu^B(n) > 0$ for some n must also induce $\mu^B(n-1) > 0$. If initial conditions at date 0 are the result of an exclusion policy that has been in place for a while, they will therefore satisfy that premise.

In summary, front-loading punishment is only one of many efficient policies when amnesty is feasible. But some of front-loading punishment is once again strictly optimal when amnesty is not an option.

II. Heterogenous default costs and exclusion length

Assume that ex-post default costs can either be low at $\tau_L = \tau - \epsilon$ or high at $\tau_H = \tau + \epsilon$ where $\epsilon > 0$ and, for concreteness, assume that these two outcomes are equally likely. The exclusion length can be set so as to dissuade both ex-post types from defaulting from strategic reasons. Instead, it can be set to dissuade just the high-default cost borrower. Finally, exclusion length can be such that it does not dissuade either borrower type from defaulting for strategic reasons.

In other words, there are three possibilities. We can set κ to solve equation (18) in the published version of our paper for $\tau = \tau_H$ in which case only low-default cost borrowers default for strategic reasons. Low-default cost agents are then excluded for the corresponding time but since they cannot be dissuaded from strategic default, it makes no sense to exclude them any longer than what is strictly necessary to keep high-cost agents in line. If this option turns out to be optimal, note that imposing a mean-preserving spread on F results in lowering the length of exclusion. Second, we can set κ to solve equation (18) for $\tau = \tau_L$ so that no agent ever defaults for strategic reasons. In that case, the mean-preserving spread results in lengthening the duration of exclusion. Third and finally, we can simply give up on dissuading any agent from strategic default by setting $\kappa = 0$.

Each of those three results of spreading F is efficient for certain parameters. This means that, in general, mean-preserving spreads on incentives to default for strategic reasons have ambiguous effects on efficient exclusion length. We can describe this ambiguity more precisely.

PROPOSITION 3: *Starting from an economy with homogenous default costs in which optimal exclusion length is positive, a mean-preserving spread in default costs raises exclusion length for ϵ small enough but must eventually drive exclusion length to zero as ϵ becomes large*

PROOF:

Start from the homogenous economy and introduce an infinitesimal spread $\tau_H - \tau_L = \epsilon > 0$. Adjusting κ by setting $\tau = \tau_L$ in equation (18) has no first order effect on any policy. Not adjusting, however, would cause half of agents with successful projects to begin defaulting for strategic reasons. Therefore adjusting by raising exclusion length infinitesimally is efficient. Once $\tau_H - \tau_L$ becomes large, high-default cost agents need not be dissuaded any longer while low default cost agents cannot be dissuaded by exclusion as τ_L becomes low and then eventually negative (these agents get positive utility from defaulting.) We now have no choice but to give up on the low-cost agents.[1] This completes the proof. ■

Local mean-preserving spreads in default costs cause exclusion length to increase because it is efficient to keep low-default cost borrowers from defaulting

[1] As the preceding discussion explained, before reaching zero there may be a point where it is optimal to only dissuade high-cost agents. Once that stage is reached, a bigger spread starts lowering exclusion length.

for strategic reasons. But as the spread in F becomes large, exclusion threats become less potent. High default-cost agents do not default anyway while very low-default cost agents simply cannot be dissuaded from doing so.