A General and Intuitive Envelope Theorem*

Andrew Clausen
University of Edinburgh

Carlo Strub
University of St. Gallen

11th January 2017

Abstract

Previous envelope theorems establish differentiability of value functions. Our techniques apply to all functions whose derivatives appear in first-order conditions. We derive first-order conditions involving the derivatives of (i) the Stackelberg follower's policy in a Stackelberg leader's problem, and (ii) a borrower's value function and default cut-off policy function in an unsecured credit economy. Our techniques also accommodate optimization problems involving discrete choices, infinite horizon stochastic dynamic programming, and Inada conditions. So we can differentiate (iii) the firm's value function in a capital adjustment problem with fixed costs, and (iv) the households' value functions in insurance arrangements with indivisible goods.

Keywords: First-order conditions, policy functions, discrete choice, Inada conditions, dynamic programming, reverse calculus.

1 Introduction

A fundamental insight of economics is that optimal choices occur where marginal benefit equals marginal cost. In simple economies, both sides of this first-order condition are exogenous, and can be assumed to exist. In recursive macroeconomies, the marginal benefit of preparing for the future is endogenous, and envelope theorems have established its existence in well-behaved convex settings. However, there are many important economic problems where it is unknown whether first-order conditions apply.

*An earlier version of this paper circulated under the title "Envelope Theorems for Non-Smooth and Non-Concave Optimization." We are very grateful to Michael Elsby, Philipp Kircher, Dirk Krueger, Tzuo Hann Law, George Mailath, Jochen Mankart, Steven Matthews, Nirav Mehta, Guido Menzio, Leonard Mirman, Kurt Mitman, John H. Moore, Georg Nöldeke, Andrew Postlewaite, Kevin Reffett, José-Victor Ríos-Rull, Felipe Saffie, József Sákovics, Ilya Segal, Shouyong Shi, Philip Short, Jonathan Thomas, Xi Weng, Tim Worrall, and Mark Wright for fruitful discussions. We especially thank an anonymous referee for suggesting we restructure our paper around the Differentiable Sandwich Lemma. Email: `andrew.clausen@ed.ac.uk` and `cs@carlostrub.ch`.

Threats to first-order conditions. *Jumps* can arise in objective functions even if all of the model primitives are continuous. For example, consider the Stackelberg duopoly game, which we study in detail in Section 2. The follower's policy function appears inside the leader's objective function. But strong conditions are required to ensure that the follower's policy function is continuous. If the follower's policy function is discontinuous, then the leader's objective function is not continuous, let alone differentiable.

Kinks can arise in objective functions even if all model primitives are differentiable. This occurs when a continuous choice is taken along side a discrete choice. For example, consider a Stackelberg leader who can build his factory in China or Europe. Assume each location has a differentiable cost function. Despite this assumption, the firm's overall cost function has a kink at the quantity where both locations are equally costly.

Hidden jumps and kinks arise when the objective is differentiable, but ingredients such as benefit and cost are not. For example, suppose that the Stackelberg leader does not know the follower's cost curve. He assigns probabilities to two different follower cost curves, and hence to two different leader benefit functions. Even if the expected benefit of selling output is differentiable, this does *not* imply that the ex post benefit curves are differentiable. They might have kinks or jumps that cancel each other out. In this case, it is impossible to write a meaningful first-order condition.

Boundary problems arise when decision makers prefer to make boundary choices, such as exhausting capacity constraints. Since first-order conditions only apply to interior solutions, economists often steer decision makers to the interior by imposing Inada conditions. This is problematic as the relevant envelope theorems depend on placing uniform bounds on derivatives. In other words, economists wishing to apply first-order conditions have an uncomfortable choice between (i) assuming Inada conditions hold to ensure that all solutions are interior, or (ii) assuming that Inada conditions do not hold, to ensure that derivatives exist.

These threats are common place in important economic problems. For example, all four threats arise in the unsecured credit market model of Arellano (2008), which we study in Section 5.1.[1] First, the borrower's future default policy appears in his objective, because it determines default risk and hence interest rates. There is no a priori reason why his policy would be differentiable. Second, the borrower has a discrete choice – whether to honour or default on debts owed – leading to kinks in his value function. Third, even if the objective is differentiable, the default policy and value function might have jumps or kinks that cancel each other out. Fourth, Arellano focuses on a utility function that satisfies the Inada conditions. These four features of Arellano's model pose difficulties to applying any of the existing envelope theorems.

[1] Below, we also discuss related work including Eaton and Gersovitz (1981), Aguiar and Gopinath (2006), Hatchondo and Martinez (2009), and Arellano and Ramanarayanan (2012).

Techniques. We devise a general and intuitive recipe that addresses these threats to first-order conditions. The recipe makes use of two new techniques. Our **Differentiable Sandwich Lemma** reformulates a classical result about subderivatives in a general and intuitive form for proving envelope theorems.[2] It establishes that a function F is differentiable at a point $\bar{c}$ if it is sandwiched between two differentiable functions U and L, as depicted in Figure 1. Specifically, the lemma applies if the two functions, which we call differentiable upper and lower support functions, satisfy (i) $U(\bar{c}) = F(\bar{c}) = L(\bar{c})$, (ii) $U(c) \geq F(c) \geq L(c)$ for all c, and (iii) L and U are differentiable at $\bar{c}$. Our lemma generalises Benveniste and Scheinkman (1979, Lemma 1) by dropping all convexity requirements.

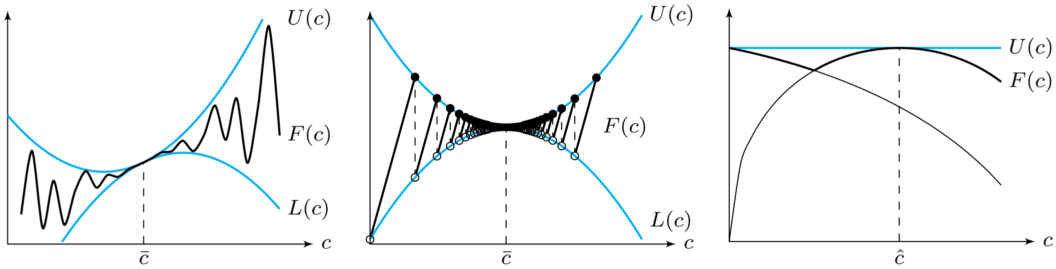

(a) F is differentiable at $\bar{c}$ (b) F is differentiable at $\bar{c}$ (c) Upper support at a max

Figure 1: Differentiable Sandwich Lemma

The lemma is well-suited to studying optimal choices. Suppose that the decision maker must make a continuous choice $c \in \mathbb{R}$ followed by a discrete choice $d \in D$. Assume that his utility function $v_d(c)$ is differentiable in c for each discrete choice d. Let $F(c) = \max_{d \in D} v_d(c)$ be the value after choosing c. Notice that at an optimal choice $(\hat{c}, \hat{d})$, the value function is sandwiched between the horizontal line $U(c) = F(\hat{c})$ and $v_{\hat{d}}$, as depicted in Figure 1c. Therefore, F is differentiable at $\hat{c}$. Milgrom and Segal (2002, Corollary 2) previously drew this conclusion for the special case that $\{v_d\}_{d \in D}$ is equi-differentiable and has uniformly bounded derivatives. Their redundant conditions conflict with Inada conditions. This means that the Differentiable Sandwich Lemma is applicable to problems with discrete choices and Inada conditions for the first time.

Our second and most novel innovation, **reverse calculus**, is the opposite of normal calculus. Whereas normal calculus establishes that $H(c) = F(c) + G(c)$ is differentiable if F and G are differentiable, reverse calculus establishes that F and G are differentiable if H is differentiable. The main requirement for reverse calculus is that each ingredient function must have an appropriate differentiable support function. For example, if

[2] Specifically, Rockafellar and Wets (1998, Proposition 8.5) establish that a function is subdifferentiable if and only if it has a differentiable lower support function. (See Clausen and Strub (2016, Appendix F) for a simpler proof.) It is then straightforward to show that a function is differentiable if and only if it is both sub- and superdifferentiable; see Kruger (2003).

$H(c) = F(c) + G(c)$, then we require F and G have differentiable lower support functions f and g at $\bar{c}$, depicted in Figure 2. Under these conditions, F is sandwiched between f and $H - g$, and is therefore differentiable.

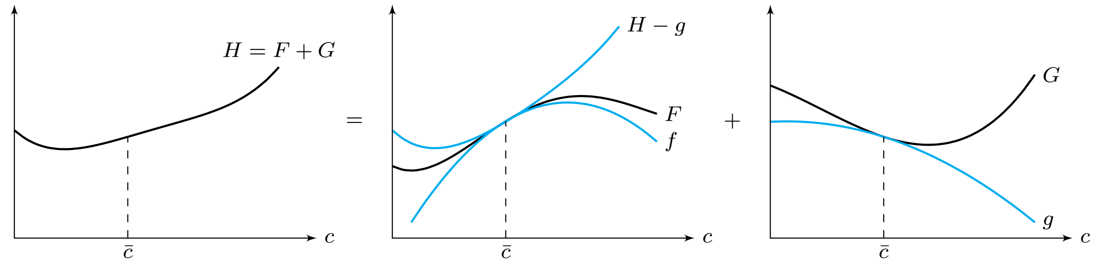

Figure 2: Reverse calculus: F is differentiable at $\bar{c}$.

Reverse calculus addresses problems whose objectives involve policy functions and/or expectations over a family of value functions. As discussed above, it is insufficient to establish that the objective function (e.g. H) is differentiable. Meaningful first-order conditions require us to establish that all of the relevant ingredient functions (e.g. F and G) are differentiable. We develop a reverse calculus for many standard operations, including addition, multiplication, function composition and upper envelopes.

Our main contribution is a **recipe for deriving first-order conditions** by using the Differentiable Sandwich Lemma and reverse calculus. We now sketch how the recipe applies to the Stackelberg duopoly problem; Section 2 gives a comprehensive analysis.

Recipe to make first-order conditions. We illustrate the recipe with a Stackelberg leader's problem. He chooses output y to maximise profit $\pi(y) = yP(y + f(y)) - C(y)$ involving the follower's endogenous best response $f(y)$, the downward-sloping inverse demand curve P and the upward-sloping production cost $C(y)$. We can assume differentiability of P and C, but not of f, and therefore not of π. The recipe proves that if $\hat{y}$ maximises π, then $f'(\hat{y})$ and $\pi'(\hat{y})$ exist, and that $\pi'(\hat{y}) = \hat{y}P'(\hat{y} + f(\hat{y}))f'(\hat{y}) + P(\hat{y} + f(\hat{y})) - C'(\hat{y}) = 0$.

Ingredients.

- 1 objective function to maximise, containing $n \geq 0$ endogenous functions.
- n differentiable support functions, one for each endogenous function, carefully chosen for step (i).

Method.

(i) Check the ingredients: you must be able to construct a **differentiable lower support function** for the objective function by substituting the support functions into the formula for the objective function. Note that the objective function might flip an upper support function into a lower support function, and vice versa. This determines whether an upper or lower support function is required for each endogenous function. (In the example, we need a differentiable upper support function F for the follower's best response f at $\hat{y}$. Replacing f with F in the formula for π gives a differentiable lower support function $L(y) = yP(y + F(y)) - C(y)$ for π at $\hat{y}$, since P is downward sloping.)

(ii) Construct the trivial **differentiable upper support function** for the objective function. (In the example, the constant function $U(y) = \pi(\hat{y})$ is a differentiable upper support function for π at $\hat{y}$.)

(iii) Use the **Differentiable Sandwich Lemma** to prove that the objective function's derivative exists and equals zero. (In the example, π is sandwiched between L and U at $\hat{y}$, so $\pi'(\hat{y})$ exists and $\pi'(\hat{y}) = U'(\hat{y}) = 0$.)

(iv) Use **reverse calculus** to prove that all derivatives in the first-order condition exist. (In the example, knowing from step (ii) that $\pi'(\hat{y})$ exists implies, by reverse calculus, that $f'(\hat{y})$ exists too.[3])

The main ingredient required for using the recipe is the collection of differentiable upper or lower support functions, one for each endogenous function.[4] How are these support functions to be found? One approach is based on "lazy decision maker" constructions that involve unreactive policy functions. Benveniste and Scheinkman's (1979) applied this strategy to construct a differentiable lower support function for value functions. This is a special case of a more general approach of constructing optimistic or pessimistic valuations. We construct support functions for optimal stopping rules based on a pessimistic option value of continuing. Similarly, we construct a support function for a Stackelberg follower policy based on pessimistic beliefs about when the follower is capacity constrained.

Contribution. The recipe allows us to provide meaningful first-order conditions that involve differentiating endogenous functions. The recipe applies even when all four threats above are present, i.e. when there are discrete choices, the primitives involve Inada conditions, and many endogenous ingredient functions are combined in expectations, budget constraints, or incentive constraints. We prove several general theorems using the re-

[3] This problem has only one endogenous function, so reverse calculus is tantamount to the implicit function theorem.

[4] In previous drafts, we proved that the recipe works if appropriate ingredients are supplied. However, this relied on an algebraic characterisation of the ingredients that lead to a differentiable lower support function for the objective. This is analogous to characterising the set of formulae that can be differentiated by the rules of calculus, which is rather tedious and unenlightening.

cipe, including a new theorem establishing first-order conditions in stochastic dynamic programming problems with discrete choices.

The recipe also allows us to resolve several open problems. We re-examine three economies in which previous authors have applied first-order conditions, where the correctness of these conditions was until now an open question. Specifically, we establish that first-order conditions do hold (i) in unsecured credit markets with endogenous default probabilities, (ii) in capital markets with fixed costs of adjustment,[5] and (iii) in informal insurance arrangements with indivisible choices.[6]

Outline. Section 2 lays out our recipe for deriving first-order conditions using a Stackelberg duopoly as a running example. Section 3 formally specifies and proves the lemmas used in the recipe. In Section 4 we compare our techniques to previous work; this includes a novel proof of the Benveniste and Scheinkman (1979) envelope theorem. In Section 5, we apply our recipe to the open questions listed above, and Section 6 concludes.

2 Illustration

This section illustrates the recipe for deriving first-order conditions through a series of examples. The examples are all Stackelberg duopoly games in which a leader's first-order condition involves the derivative of a follower's policy function. Other related problems are explored by Kydland and Prescott (1977) and Ljungqvist and Sargent (2012, Chapter 19). The textbook analysis requires strong convexity and twice-differentiability assumptions, which we relax by introducing capacity constraints. The follower's policy has a kink where he switches from being constrained to unconstrained. We establish that the leader steers the follower away from these kinks.

The first example illustrates how to apply the Differentiable Sandwich Lemma to obtain first-order conditions when there is only one policy function. In the second example, there are two policy functions. We apply reverse calculus to establish that the follower steers the leader away from the kinks in both policy functions. While reverse calculus applies generally to optimisation problems, the third example shows that it generically fails elsewhere.

Textbook Stackelberg Competition. To fix notation, we review the textbook analysis (e.g. Varian (1992, Section 16.6)) of a Stackelberg duopoly. A leader and a follower choose their output levels, y_1 and y_2 sequentially, which costs them $C(y_1)$ and $C(y_2)$. The output

[5] Below, we discuss the work of Harrison, Sellke and Taylor (1983), Caballero and Engel (1999), Cooper and Haltiwanger (2006), Gertler and Leahy (2008), Khan and Thomas (2008b) and Elsby and Michaels (2014).

[6] Below, we discuss the work of Thomas and Worrall (1988, 1990), Kocherlakota (1996), Ligon, Thomas and Worrall (2002), Koeppl (2006), Rincón-Zapatero and Santos (2009), and Morten (2015).

is sold at the market price, $P(y_1 + y_2)$. Firm i earns profits $\pi_i(y_1, y_2) = y_i P(y_1 + y_2) - C(y_i)$. The follower chooses $y_2 = f(y_1)$ by solving

$$f(y_1) = \arg\max_{y_2} \pi_2(y_1, y_2) \tag{1}$$

and the leader chooses the y_1 that maximises his objective

$$\phi_1(y_1) = \pi_1(y_1, f(y_1)). \tag{2}$$

The textbook first-order conditions for the follower and leader are

$$P(y_1 + y_2) + P'(y_1 + y_2) y_2 = C'(y_2) \text{ and} \tag{3}$$

$$P(y_1 + f(y_1)) + P'(y_1 + f(y_1))(1 + f'(y_1)) y_1 = C'(y_1). \tag{4}$$

The derivatives of P, f, and C appear in (4). The demand function P and cost function C are exogenous, so we can assume they are differentiable. However, we can not assume that the follower chooses a differentiable policy f.

The textbook solution is to assume the cost and demand functions are strictly convex/concave and twice differentiable. Under these assumptions, (3) implicitly defines the follower's policy function, which is differentiable by the implicit function theorem. Therefore, the leader's first-order condition (4) holds at her optimal choices.

Example 1: Stackelberg with Capacity Constraints. Suppose the follower has a capacity constraint Y. The follower's cost function is

$$C(y) = \begin{cases} c(y) & \text{if } y \leq Y, \\ \infty & \text{if } y > Y, \end{cases} \tag{5}$$

where $c(\cdot)$ is twice differentiable and strictly convex.

The follower's best response has an upward kink, depicted in Figure 3a, where his capacity constraint transitions from binding to non-binding. This translates into a downward kink in the leader's objective, depicted in Figure 3b. Would the leader ever choose this kink, thus invalidating the first-order condition (4)?

The answer is no. To prove this, we will construct a differentiable sandwich around the leader's objective $\phi_1(\cdot)$ at the optimal choice $\hat{y}_1$. First, we construct the bottom half of the sandwich, corresponding to step (i) of the recipe. Since the follower's policy (depicted in Figure 3a) is the lower envelope of two differentiable functions, it has a differentiable upper support function $F(\cdot)$ at $\hat{y}_1$. The follower's policy enters the leader's objective through the downward sloping demand curve P. This means that $L(y_1) = \pi_1(y_1, F(y_1))$ is a lower support function for the leader's objective at $\hat{y}_1$.

Second, we construct the top half of the sandwich as the constant function $U(y_1) = \phi(\hat{y}_1)$. This corresponds to step (ii) of the recipe.

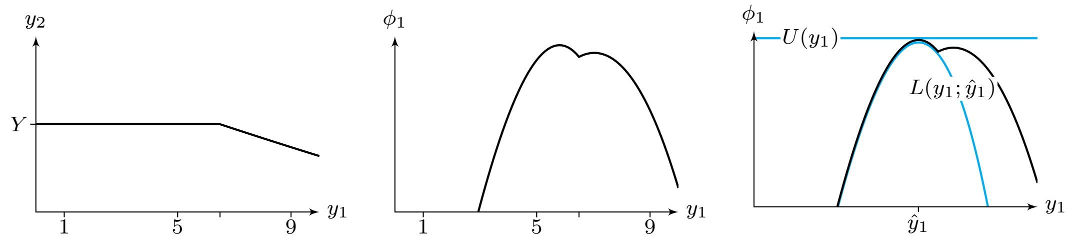

(a) Follower's policy (b) Leader's profit (c) Differentiable sandwich

Figure 3: Stackelberg with a capacity constrained follower

These support functions form a sandwich, illustrated in Figure 3c. By the Differentiable Sandwich Lemma, the leader's objective is differentiable at $\hat{y}_1$, where it satisfies the first-order condition

$$\phi_1'(\hat{y}_1) = U'(\hat{y}_1) = 0. \tag{6}$$

This completes step (iii) of the recipe.

However, we have not yet established (4), which is a more useful first-order condition. In particular, we have not yet determined whether f is differentiable at $\hat{y}_1$. This is a reverse calculus problem: we have established that the left side of (2) is differentiable, and we would now like to infer that the policy f on the right side is differentiable. In this example, there is a simple solution. Since (2) implicitly defines f near $\hat{y}_1$, the implicit function theorem implies f is differentiable at $\hat{y}_1$. We conclude that the leader's first-order condition (4) holds at optimal choices. This completes step (iv) of the recipe.

This example illustrates how the Differentiable Sandwich Lemma and a simple form of reverse calculus can be applied to establish first-order conditions. The main task was constructing the bottom half of the sandwich. We constructed a differentiable upper support function for the follower's policy, which we used to construct a differentiable lower support function for the leader's objective. Since there was only one endogenous function, the reverse calculus step only required the implicit function theorem.

Example 2: Stochastic Stackelberg with capacity constraints. The previous example only required a simple reverse calculus step, because there was only one policy function. We now consider a problem with two policy functions. We extend Example 1 by assuming that the follower has privately known costs. Specifically, the leader assigns probabilities p_A and p_B to the follower's production cost being $C_A(\cdot)$ or $C_B(\cdot)$, respectively. The follower now has two policies f_A and f_B, one for each cost function. The leader's problem is to choose output y_1 to maximise her expected profit

$$\phi_1(y_1) = \sum_{z \in \{A,B\}} p_z y_1 P(y_1 + f_z(y_1)) - C(y_1). \tag{7}$$

We would like to determine whether her first-order condition

$$\sum_{z\in\{A,B\}} p_z \left\{ P(y_1 + f_z(y_1)) + y_1 P'(y_1 + f_z(y_1))(1 + f_z'(y_1) \right\} - C'(y_1) = 0 \qquad (8)$$

holds at her optimal choice $\hat{y}_1$. As before, the follower's policy functions have one kink each, depicted in Figure 4a. The corresponding leader objective function (depicted in Figure 4b) inherits both kinks, but neither kink is an optimal choice. Does the leader always steer the follower away from the kinks in his policy functions?

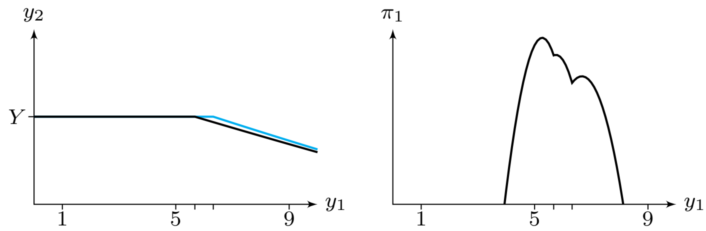

(a) Follower's policies (b) Leader's profit

Figure 4: Stackelberg follower with privately known cost

As in the previous example, we can construct a differentiable upper support function $F_z(\cdot)$ for each policy $f_z(\cdot)$ at $\hat{y}_1$. Steps (i) to (iii) of the recipe follow a similar course to establish that the first-order condition $\phi_1'(\hat{y}_1) = 0$ holds.

We show that the follower's policy functions $f_z(\cdot)$ are differentiable at $\hat{y}_1$. In the previous example, this was a straightforward application of the implicit function theorem. However, one equation can not implicitly define two policy functions. Instead, we apply our reverse calculus summation rule to

$$\sum_{z\in\{A,B\}} p_z y_1 P(y_1 + f_z(y_1)).$$

Since this sum is differentiable, and each term has a differentiable lower support function $L_z(y_1) = p_z y_1 P(y_1 + F_z(y_1))$, the rule implies that each term is differentiable at $\hat{y}_1$. Therefore, both policy functions are differentiable at optimal choices, and the first-order condition (8) holds. This completes step (iv) of the recipe.

When there are two or more policy functions, the implicit function theorem can not be applied to establish they are differentiable. This example showed how reverse calculus applies in these situations. In fact, we did not need to impose any additional conditions. This is not a coincidence; if the first step of the recipe succeeds, then the other steps will also succeed.

Example 3: Stackelberg leader's value. In the first two examples, we applied reverse calculus to show that the leader chooses differentiable points of the follower's policies. In this example, we illustrate how reverse calculus might fail. We focus our attention on the follower's value function,[7]

$$V_2(y_1) = f(y_1)P(y_1 + f(y_1)) - C(f(y_1)). \tag{9}$$

It is straightforward to show that V_2 is a concave function. Therefore, the Benveniste and Scheinkman (1979) envelope theorem (which we prove in Section 4.1) implies that V_2 is globally differentiable. It is tempting to deduce that the follower's policy, f, which appears twice on the right side of (9), must also be globally differentiable. However, in Figure 3a, we saw that f is *not* globally differentiable!

Why does reverse calculus not apply here? The right side of (9) is the difference between revenue and cost, which have identical kinks that cancel each other out. This cancellation is depicted in Figure 5. The relevant reverse calculus rule is the implicit function theorem, because the policy function $y_2 = f(y_1)$ is implicitly defined by the equation $\psi(y_1, y_2) \equiv V_2(y_1) - y_2 P(y_1 + y_2) + C(y_2) = 0$. The implicit function theorem requires that $\frac{\partial}{\partial y_2}\psi(y_1, y_2)$ be non-zero at $(y_1, f(y_1))$. But it is *always* zero – this is the follower's first-order condition for choosing y_2.

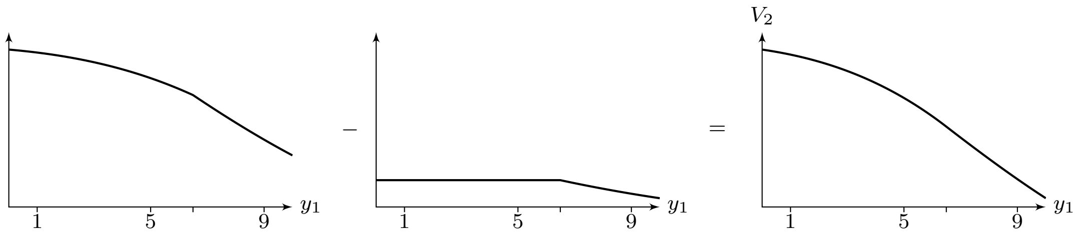

(a) Follower's optimal revenue, $f(y_1)P(y_1 + f(y_1))$

(b) Follower's optimal cost, $C(f(y_1))$

(c) Follower's value function

Figure 5: Kinks exactly cancel, endogenously

Kink cancellation is not special to value functions in Stackelberg games. It is a property of generic concave dynamic programming problems of the form

$$\begin{aligned} V(a) &= \max_{a'} u(a, a') + \beta V(a') & (10) \\ &= u(a, f(a)) + \beta V(f(a)), & (11) \end{aligned}$$

where u is concave and once (but not twice) differentiable, and f is the policy function. Again, V is globally once (but not twice) differentiable. Santos (1991) showed that f is only differentiable where V is twice differentiable. We conclude that the left side of (11) is differentiable, but that the two terms on the right side are not; they contain kinks that cancel each other out where V is not twice differentiable.

[7] Value functions of other players arise in problems with promise-keeping constraints.

This is a cautionary example. When the conditions of reverse calculus are not met, kinks in ingredient functions might be hidden by kink cancellation. This occurs in generic concave dynamic programming problems.

3 Envelope Recipe Techniques

The key logic of the recipe is contained in the Differentiable Sandwich Lemma and reverse calculus. We presents proofs in full generality of these techniques.

3.1 Differentiable Sandwich Lemma

Before stating the lemma, we need to be precise about what a derivative is. Since we would like to accommodate many continuous choices (such as asset portfolio choices), we use the standard multidimensional definition of differentiability. This definition ensures that the chain rule and other calculus identities are valid.

Definition 1. A function $F : C \to \mathbb{R}$ with domain $C \subseteq \mathbb{R}^n$ is **differentiable** at $c \in \operatorname{int}(C)$ if there is some row vector m with $m^\top \in \mathbb{R}^n$ such that

$$\lim_{\Delta c \to 0} \frac{F(c + \Delta c) - F(c) - m\,\Delta c}{\|\Delta c\|} = 0. \tag{12}$$

Such an m is the **derivative** of F at c, and is denoted $F'(c)$.

In fact, this definition is almost identical to the case where the domain is a subset of a Banach space $(X, \|\cdot\|)$, and our results generalize without amendment.[8]

Lemma 1 (Differentiable Sandwich Lemma)**.** *If F is differentiably sandwiched between L and U at $\bar{c}$ then F is differentiable at $\bar{c}$ with $F'(\bar{c}) = L'(\bar{c}) = U'(\bar{c})$.*

Proof. The difference function $d(c) = U(c) - L(c)$ is minimized at $\bar{c}$. Therefore, $d'(\bar{c}) = 0$ and we conclude $L'(\bar{c}) = U'(\bar{c})$.

Let $m = L'(\bar{c}) = U'(\bar{c})$. For all Δc,

$$\begin{aligned}\frac{L(\bar{c} + \Delta c) - F(\bar{c}) - m\,\Delta c}{\|\Delta c\|}& \\ \leq \frac{F(\bar{c} + \Delta c) - F(\bar{c}) - m\,\Delta c}{\|\Delta c\|} &\leq \frac{U(\bar{c} + \Delta c) - F(\bar{c}) - m\,\Delta c}{\|\Delta c\|}.\end{aligned} \tag{13}$$

[8] In Banach spaces, the derivative m is called a "Fréchet derivative" and lies in the topological dual space $X^* = \{m : X \to \mathbb{R} \text{ such that } m \text{ is linear and continuous}\}$. For our purposes, it is unnecessary to define a topology on X^* because all limits are taken in $(X, \|\cdot\|)$ and $\mathbb{R}$.

Consider the limits as $\Delta c \to 0$. Since $L'(\bar{c}) = U'(\bar{c}) = m$, the limits of the first and last fractions are 0. By Gauss' Squeeze Theorem, we conclude that the limit in the middle is also 0, and hence that F is differentiable at $\bar{c}$ with $F'(\bar{c}) = m$. □

Remark 3.1. The Differentiable Sandwich Lemma also applies when $F : C \to \mathbb{R}$ is sandwiched between L and U on an open neighbourhood of $\bar{c}$.

3.2 Maximum Lemma

Step (ii) of the recipe requires us to construct a differentiable upper support function of the objective. This problem has a trivial solution: a horizontal line (or hyperplane) through the maximum of the objective (see Figure 3c).

Lemma 2 (Maximum Lemma). *Let $\phi : C \to \mathbb{R}$ be a function. If $\hat{c} \in \text{int}(C)$ maximises ϕ, then $U(c) = \phi(\hat{c})$ is a differentiable upper support function of ϕ.*

3.3 Reverse Calculus

Calculus involves rules such as "if F and G are differentiable at $\bar{c}$, then $H(c) = F(c) + G(c)$ is also differentiable at $\bar{c}$." Reverse calculus rules go in the opposite direction. We provide the most important rules here, and additional rules for convex combinations and endogenous function composition in Appendix A. We also omit the corresponding rules for subtraction and division, which would involve both upper and lower support functions.

Lemma 3 (Reverse Calculus). *Suppose $F : C \to \mathbb{R}$, and $G : C \to \mathbb{R}$ have differentiable lower support functions f, and g respectively at $\bar{c}$.*

(i) If $H(c) = F(c) + G(c)$ is differentiable at $\bar{c}$, then F is differentiable at $\bar{c}$.

(ii) If $H(c) = F(c)G(c)$ is differentiable at $\bar{c}$ and $F(\bar{c}) > 0$ and $G(\bar{c}) > 0$, then F is differentiable at $\bar{c}$.

(iii) If $H(c) = \max\{F(c), G(c)\}$ is differentiable at $\bar{c}$ and $F(\bar{c}) = H(\bar{c})$, then F is differentiable at $\bar{c}$.

(iv) If $H(c) = J(F(c))$ and $J : \mathbb{R} \to \mathbb{R}$ are differentiable at $\bar{c}$ and $F(\bar{c})$ respectively with $J'(F(\bar{c})) \neq 0$, then F is differentiable at $\bar{c}$.

Proof. Let f and g be differentiable lower support functions of F and G at $\bar{c}$. For (i)–(iii), we sandwich F between f and an appropriate differentiable upper support function U and apply the Differentiable Sandwich Lemma (Lemma 1). Appropriate upper support functions are (i) $U(c) = H(c) - g(c)$, (ii) $U(c) = H(c)/g(c)$, and (iii) $U(c) = H(c)$.

For (iv), $F(c) = J^{-1}(H(c))$ is differentiable at $\bar{c}$ by the inverse function theorem and the chain rule. □

4 Special Cases and Related Literature

We apply our recipe to various classes of decision problems, and contrast our results to previous literature. The main advances are that we can study all types of endogenous functions (not just value functions), we accommodate non-smooth problems involving discrete choices or boundaries, and we accommodate Inada conditions that are often imposed to simplify first-order conditions.

4.1 Value Functions in Smooth Concave Problems

Benveniste and Scheinkman (1979) study *value* functions in smooth concave dynamic programming problems, but not *policy* functions. Their main theorem establishes that value functions in this setting are differentiable. The Differentiable Sandwich Lemma leads to an elementary proof of their theorem.

Problem 1. Consider the following dynamic programming problem:

$$V(c) = \sup_{c' \in \{\bar{c}' : (c, \bar{c}') \in \Gamma\}} u(c, c') + \beta V(c'), \tag{14}$$

where the domain of V is C. We assume that (i) Γ is a convex subset of $C \times C$, (ii) u is concave, and (iii) $u(\cdot, c')$ and $u(c, \cdot)$ are differentiable, respectively.

Theorem 1 (Benveniste-Scheinkman Theorem). *If $\hat{c}'$ is an optimal choice at state $c \in \operatorname{int}(\{\bar{c} : (\bar{c}, \hat{c}') \in \Gamma\})$, then V is differentiable at c with $V_1(c) = u_1(c, \hat{c}')$.*

Proof. V is concave because u is concave and Γ is convex. Hence, the supporting hyperplane theorem can be applied to the hypograph of V to construct a linear upper support function U that touches V at c. We construct the differentiable lower support function $L(c) = u(c, \hat{c}') + \beta V(\hat{c}')$. Lemma 1 delivers the conclusions. $\square$

Graduate economics textbooks such as Stokey and Lucas (1989) do not provide a self-contained proof of this theorem. Our proof is short and elementary, and therefore suitable for junior graduate students. The original proof is based on a sandwich lemma, which Benveniste and Scheinkman (1979) prove with the help of Rockafellar (1970, Theorem 25.1). (Their lemma imposes a redundant assumption, that the lower support function be concave.) Mirman and Zilcha (1975, Lemma 1) prove a one-dimensional special case using Dini derivatives rather than sandwiches.

4.2 Value Functions in Non-Smooth Problems

Milgrom and Segal (2002) study the differentiability of *value* functions and objective functions without making any topological or convexity assumptions. Our main contribution is that we also study *policy* functions. In addition, we present two generalisations of their theorems. Our first generalisation accommodates Inada conditions, which are frequently employed to ensure first-order conditions apply. We then generalise further to accommodate stochastic dynamic programming problems.

Their envelope theorems are the first to accommodate discrete choices. Specifically, they consider value functions of the form $\phi(c) = \sup_d f(c, d)$, where $\{f(\cdot, d)\}_{d \in D}$ is an arbitrary collection of differentiable functions. Here, c and d represent continuous and discrete choices, such as a quantity and factory location choice. Their Corollary 2 establishes that discrete choices only lead to downward kinks in ϕ, which decision makers would always avoid. Their result is a special case of the following theorem:

Theorem 2. *If* $(\hat{c}, \hat{d}) \in \arg\max f$ *then* ϕ *is differentiable at* $\hat{c}$ *with* $\phi'(\hat{c}) = f_1(\hat{c}, \hat{d}) = 0$.

Proof. The objective ϕ is differentiably sandwiched between $L(c) = f(c, \hat{d})$ and $U(c) = f(\hat{c}, \hat{d})$ at $\hat{c}$. □

We illustrate this theorem with an example, before comparing Milgrom and Segal's work with our own.

Example 4: Stackelberg with Factory Choice. Suppose the follower has access to one of two factories, which have cost functions C^1 and C^2, respectively (which are twice differentiable and strictly convex as usual). The follower's overall cost function is therefore $C(y) = \min\{C^1(y), C^2(y)\}$. The follower's profit function has a downward kink at $\tilde{y}_2$ where both factories are equally costly, as depicted in Figure 6a.

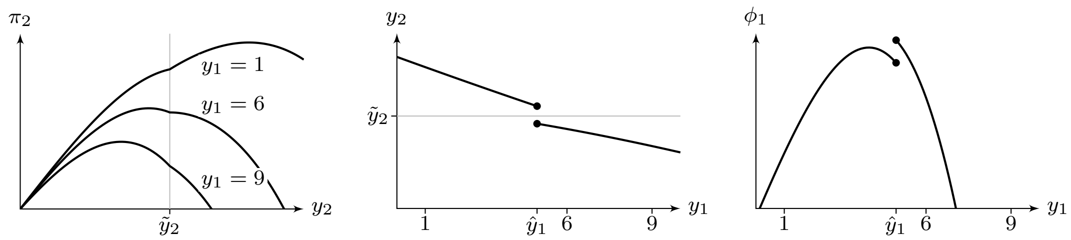

(a) Follower's profit, $\pi_2(y_1, y_2)$ (b) Follower's policy, $f(y_1)$ (c) Leader's objective function, $\phi_1(y_1) = \pi_1(y_1, f(y_1))$

Figure 6: Stackelberg when the follower has a discrete choice

Theorem 2 establishes that decision makers never choose kinks like $\tilde{y}_2$, i.e. that arise from discrete choices. This means that decision makers only choose points where their objective is differentiable. However, this does not mean that the textbook analysis of Stackelberg games generalises to this example. The follower's policy function and leader's value function, depicted in Figure 6b and Figure 6c, are discontinuous. The leader's first-order condition (4) does not hold at her optimal choice, $\hat{y}_1$.

Milgrom and Segal (2002, Corollary 2) is a special case of our Theorem 2. Their version imposes redundant equidifferentiability and bounded derivative conditions. These extra conditions ensure that directional derivatives of ϕ exist globally, which their proof makes use of. Similarly, other papers make assumptions including Lipschitz continuity,[9] and supermodularity[10] to ensure the existence of directional derivatives. Our method does not require any such assumptions.

Both of their redundant conditions are problematic. The uniformly bounded derivative condition conflicts with Inada conditions. Inada conditions are imposed to ensure first-order conditions hold by directing optimal choices away from boundaries.

The equidifferentiability condition is problematic in dynamic problems, when a discrete decision is taken in every period. The number of combinations of discrete choices increases exponentially as the number of periods increases, which can lead the number of kinks to grow rapidly. In other words, the kinks from tomorrow's value function propagate into today's value function, as depicted in Figure 7a. In infinite horizon problems, the set of combinations of discrete choices is (uncountably) infinite. This can cause directional differentiability of value functions can fail. For example, the "bouncing ball" function depicted in Figure 7b has no directional derivatives at $c = 0$, even though it is the upper envelope of functions with uniformly bounded derivatives.[11]

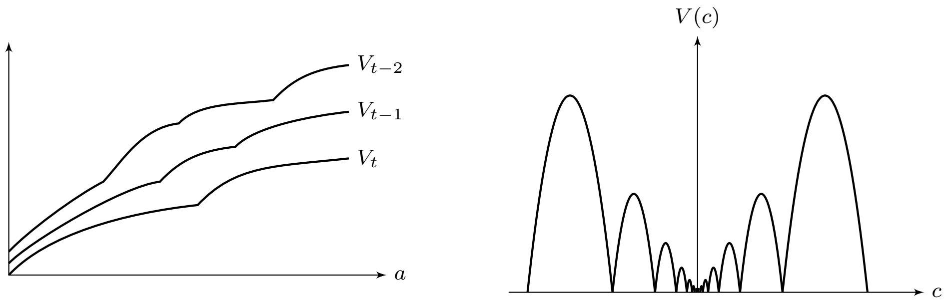

(a) Backward propagation of downward kinks

(b) The bouncing ball function has no directional derivatives at $c = 0$

Figure 7

Our differentiable sandwich approach overcomes these obstacles, so we can provide the first general envelope theorem for dynamic programming problems involving discrete choices.

[9] Clarke (1975)

[10] Amir, Mirman and Perkins (1991)

[11] The bouncing ball function is the upper envelope of a set of parabolas, $\{v(\cdot, d)\}_{d \in D}$ where $v(c, d) = -\frac{2}{|d|}(c-d)(c-\frac{1}{2}d)$ and $D = \{s2^{-n} : s \in \{-1, 1\}, n \in \mathbb{N}\}$. On the relevant parts of the domains, their derivatives lie in $[-1, 1]$.

Problem 2. Consider the following stochastic dynamic programming problem:

$$V(c,d,\theta) = \sup_{c',d'} u(c,c';d,d';\theta) + \beta \sum_{\theta' \in \Theta} \pi_{\theta\theta'} V(c',d',\theta'),$$
$$\text{s.t. } (c,c';d,d';\theta) \in \Gamma,$$

where the domain of V is $\Omega \times \Theta$. We assume that $u(\cdot, c'; d, d'; \theta)$ and $u(c, \cdot; d, d'; \theta)$ are differentiable.

Suppose that $C(c,d,\theta)$ and $D(c,d,\theta)$ are optimal choices at the state (c,d,θ).

Definition 2. The set of feasible one-shot deviations from the optimal policies at state (c,d,θ) is

$$\Lambda(c,d,\theta) = \left\{ c' : (c,c';d,\bar{d}';\theta) \in \Gamma, \text{ and for all } \theta', (c', \bar{c}''(\theta'); d', \bar{d}''(\theta'); \theta') \in \Gamma \right\},$$

where $(\bar{c}', \bar{d}') = (C(c,d,\theta), D(c,d,\theta))$, and $(\bar{c}''(\theta'), \bar{d}'(\theta')) = (C(\bar{c}', \bar{d}', \theta'), D(\bar{c}', \bar{d}', \theta'))$.

Theorem 3. *Let $(\hat{c}', \hat{d}') = (C(c,d,\theta), D(c,d,\theta))$ be optimal choices at state (c,d,θ). If $\hat{c}'$ is an interior choice, i.e. $\hat{c}' \in \text{int}(\Lambda(c,d,\theta))$, then (i) $V(\cdot, \hat{d}')$ is differentiable at $\hat{c}'$ and (ii) $\hat{c}'$ satisfies the first-order condition*

$$-u_{c'}(c,\hat{c}';d,\hat{d}';\theta) = \beta \sum_{\theta'} \pi_{\theta\theta'} V_c(\hat{c}', \hat{d}', \theta') = \beta \sum_{\theta'} \pi_{\theta\theta'} u_c(\hat{c}', \hat{c}''(\theta'); \hat{d}', \hat{d}''(\theta'); \theta'),$$

where $(\hat{c}''(\theta'), \hat{d}'(\theta'))$ are shorthand for $(C(\hat{c}', \hat{d}', \theta'), D(\hat{c}', \hat{d}', \theta'))$.

Proof. We assumed that $\hat{c}'$ maximises

$$\phi(c') = u(c,c';d,\hat{d}';\theta) + \beta \sum_{\theta' \in \Theta} \pi_{\theta\theta'} V(c', \hat{d}', \theta'), \tag{15}$$

where the domain of ϕ is $\Lambda(c,d,\theta)$. The proof follows the four steps of the recipe. First, the value function V has a differentiable lower support function at $(\hat{c}', \hat{d}', \theta')$,

$$L(c'; \hat{c}', \hat{d}', \theta') = u(c', \hat{c}''(\theta'); d', \hat{d}''(\theta'), \theta') + \beta \sum_{\theta'' \in \Theta} \pi_{\theta'\theta''} V(\hat{c}''(\theta'), \hat{d}''(\theta'), \theta''). \tag{16}$$

Therefore, ϕ is sandwiched between a corresponding differentiable lower support function, and a constant upper support function at $\hat{c}'$m completing the second step. Third, the Differentiable Sandwich Lemma implies ϕ is differentiable at $\hat{c}'$. Fourth, the addition rule of reverse calculus implies V is differentiable at each $(\hat{c}', \hat{d}', \theta')$. □

4.3 Policy Functions in Smooth Concave Problems

Previous theorems about the differentiability of policy functions focus on smooth deterministic concave settings. Our main contribution is that our recipe can be applied without any of these conditions. For example, in our Stackelberg illustration in Example 1, we accommodate boundaries and discontinuous marginal utilities. In option value problems (such as the unsecured credit application in Section 5.1), we accommodate discrete choices and shocks.

Araújo and Scheinkman (1977) study the differentiability of policy functions, further restricting attention to deterministic one-good problems in which the utility function is twice differentiable. Their proof is based on generalising the implicit function theorem to infinite-dimensional spaces. Thus, their approach is similar to our use of the implicit function theorem in Example 1. Reverse calculus is necessary to extend this approach beyond one endogenous function.

Santos (1991) drops the one-good restriction, but retains the concavity and twice-differentiability conditions. As mentioned above, he begins by noticing that differentiability of policy functions is equivalent to twice-differentiability of value functions. He proceeds to construct a convergent sequence of quadratic approximations of the value function to establish this twice-differentiability. Like Araújo and Scheinkman (1977), his results only address the differentiability of policy functions and value functions. He does not accommodate qualitatively different endogenous functions such as the cut-off policies in Section 5.1.

5 Applications

5.1 Unsecured Credit

Our first application is about unsecured debt contracts where borrowers may decide to either repay in full or to default. We focus on markets without collateral such as sovereign debt markets. The punishment for default is exclusion from the credit market thereafter. Nevertheless, default occasionally occurs so interest paid by the borrower must compensate for the default risk.[12] For this reason, the interest charged is non-linear and determined by a recursive relationship with the borrower's value function. If the interest rates are low, then the borrower's value of honouring debt contracts is high because rolling over debt is cheap. Conversely, if the borrower's value of repaying is high tomorrow, then the default risk today is low. This recursive relationship determines interest rates as a function of loan sizes.

All four threats to first-order conditions from the introduction are present. First, the borrower's future default policy appears in his objective, because it determines default

[12] Default need not be inefficient compared to risk-free debt, as it implements risk-sharing.

risk and hence interest rates. There is no a priori reason why his policy would be differentiable. Second, the borrower has a discrete choice – whether to honour or default on debts owed – leading to kinks in his value function. Third, even if the objective is differentiable, the default policy and value function might have jumps or kinks that cancel each other out. Fourth, we impose (problematic) Inada conditions so that the borrower makes interior consumption choices.

We apply our recipe to establish that both endogenous functions – the value function and the interest rate – are differentiable at optimal debt choices (except at the risk-free credit limit). Hence, we derive a first-order condition involving a marginal interest rate and a marginal continuation value. We also apply reverse calculus to express the marginal interest rate in terms of the cut-off policy; this allows us to express the first-order condition in terms of quantities rather than prices. Finally, we apply our recipe to characterise the borrower's credit limit and establish that this limit is never exhausted.

We build on the unsecured credit analysis by Arellano (2008) which is in the tradition of Eaton and Gersovitz (1981). Arellano carefully analyses it theoretically and numerically. She also sketches a Laffer curve for the debt choice, but – without first-order conditions – does not characterise borrower behaviour along it. The following three papers apply some Euler equations, with the first explicitly acknowledging that they lack justification for differentiating the interest rates with respect to loan size. We provide a justification. Aguiar and Gopinath (2006) dropped a detailed discussion of their heuristic Euler equation from their NBER working paper version. Similarly, Arellano and Ramanarayanan (2012) use heuristic Euler equations to compare maturity structures of loans. Finally, Hatchondo and Martinez (2009) discuss an Euler equation, implicitly assuming differentiability of interest rates.

Model. A risk-averse borrower has a differentiable utility function u and discount factor $\beta \in (0, 1)$. The borrower's marginal value of consumption at zero is infinite, i.e. $\lim_{c \to 0^+} u_1(c) = \infty$. Every period, the borrower receives an endowment x which is independently and identically distributed with density $f(\cdot)$ on the support $[x^{\min}, x^{\max}]$. We assume the borrower's endowment is bounded away from zero, i.e. $x^{\min} > 0$. To smooth out endowment shocks, the borrower may take out loans from a lender with deep pockets. We focus our attention on debt contracts of the following form. The borrower promises to pay a lender b' in the following period, although both understand that the borrower only has an incentive to honour the promise if tomorrow's x' lies in some set H'. Thus, a debt contract consists of (b', H'). The lender is risk-neutral, discounts time at the same rate, and is therefore willing to pay $\beta \int_{H'} f(x')\, dx' b'$ in return for the promise. If the borrower defaults – regardless of whether $x' \in H'$ – he is excluded from credit markets thereafter. We also accommodate an additional exogenous sanction of $s \geq 0$ units of consumption every period for defaulting, which reflects the difficulty of settling non-financial transactions without credit.[13] The borrower's autarky value after defaulting is

$$W_{\text{aut}}(x) = u(x - s) + \beta \int_{[x^{\min}, x^{\max}]} W_{\text{aut}}(x') f(x') dx'. \tag{17}$$

The lender only agrees to the contract (b', H') if the borrower has an incentive to honour the promise for the proposed endowments H'. Specifically, the borrower's value of repaying b' at an honour endowment $x' \in H'$, denoted $W_{\text{hon}}(b', x')$, should not be less than the autarky value $W_{\text{aut}}(x)$. The borrower's value of honouring debts is therefore[14]

$$\begin{aligned} W_{\text{hon}}(b, x) = \max_{c, b', H'} \; & u(c) + \int_{[x^{\min}, x^{\max}]} \max\{W_{\text{aut}}(x'), W_{\text{hon}}(b', x')\} f(x') dx', \\ \text{s.t.} \quad & c + b = x + \left[\beta \int_{H'} f(x')\, dx'\right] b', \\ & W_{\text{hon}}(b', x') \geq W_{\text{aut}}(x') \text{ for all } x' \in H', \\ & b' \leq b^{\text{ponzi}}. \end{aligned} \tag{18}$$

The last constraint rules out Ponzi schemes and the b^{ponzi} parameter may be arbitrarily large.

Reformulation. We reformulate this problem by making two simplifications. First, Arellano (2008, Proposition 3) established that because x is IID, the honour set H' chosen by the borrower is determined by a cut-off rule $y(\cdot)$ so that the borrower honours his debt at state (b', x') if and only if $x' \geq y(b')$. In other words, the borrower only ever chooses debt contracts of the form $(b', H') = (b', [y(b'), x^{\max}])$, so debt contracts are characterised by b' alone. This means we may denote the price of debt $q(b')$ as a function of b'. Second, we substitute the budget constraint into the objective, so that the borrower's only choice is his future debt obligation b'. The reformulated problem becomes

$$W_{\text{hon}}(b, x) = \max_{b' \leq b^{\text{ponzi}}} u(x + q(b')b' - b) + \beta W(b'), \tag{19}$$

where

$$W(b') = \int_{[x^{\min}, x^{\max}]} \max\{W_{\text{aut}}(x'), W_{\text{hon}}(b', x')\} f(x') dx', \tag{20a}$$

$$q(b') = \beta[1 - F(y(b'))], \tag{20b}$$

$$y(b') = \min\left(\left\{x' \in [x^{\min}, x^{\max}] : W_{\text{hon}}(b', x') \geq W_{\text{aut}}(x')\right\} \cup \{x^{\max}\}\right). \tag{20c}$$

[13] Exogenous sanctions are often included in unsecured credit models, so we include them to show the generality of our technique. Without them, Bulow and Rogoff (1989) show that exclusion from credit markets alone is an insufficient punishment for enforcing debt contracts if the borrower can make private investments.

[14] We mention some technicalities: (i) the borrower should be constrained to choosing a measurable honour set, and (ii) since we focus on first-order conditions, we take it for granted that the value function in the sequence problem is the unique solution to the Bellman equation.

We denote optimal policy functions by $\hat{b}'(b,x)$.[15]

The objective (19) has two endogenous functions, q and W, which we will show are not globally differentiable. The value function W has downward kinks at states of indifference between honouring and defaulting, as in the Stackelberg factory location choice problem depicted in Figure 6a. Similarly, we have no a priori knowledge of the differentiability of the debt price q.

We will follow the four steps of the recipe to establish that at optimal choices, first-order conditions hold and that q and W are differentiable. However, we find that there is one exception: the debt price exhibits an upward kink at the risk-free credit limit. This means that first-order conditions are inapplicable when the borrower chooses to exhaust his risk-free credit limit.

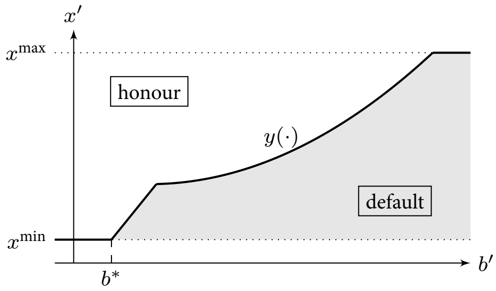

(a) Borrowers default when $x' < y(b')$

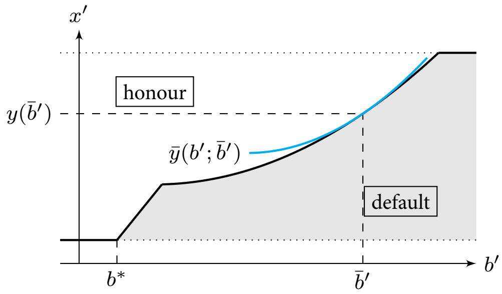

(b) A "pessimistic" borrower undervalues honouring debts, and defaults too much

Figure 8: The default cut-off rule

Ingredients: Differentiable Lower Support Functions. The most important ingredient of the recipe is providing appropriate differentiable support functions for the endogenous functions. The problem of constructing a differentiable lower support function for the debt price $q(\cdot)$ is equivalent to that of constructing a differentiable upper support function for the cut-off rule $y(\cdot)$, illustrated in Figure 8a. For debts below some threshold b^*, the borrower always honours his obligations, so the cut-off $y(\cdot)$ is constant and hence differentiable on $[-\infty, b^*)$. At each debt level $\bar{b}' > b^*$, we now construct a differentiable upper support function for $y(\cdot)$. We consider a pessimistic borrower that undervalues honouring debts, and hence uses a higher cut-off than $y(\cdot)$. Specifically we consider a pessimistic borrower who incorrectly anticipates the state to be $(b', x') = (\bar{b}', y(\bar{b}'))$, i.e. he anticipates his state will be on the cut-off. In unanticipated states, he chooses his debt to be $\hat{b}''(\bar{b}', y(\bar{b}'))$ independently of the realized endowment x'. His consumption is adjusted by the differences from the anticipated endowment and debt. This pessimistic

[15] The borrower might be indifferent between several optimal policies.

borrower's value function is

$$L(b', x'; \bar{b}') = u(x' - b' + q(\bar{b}'')\bar{b}'') + \beta W(\bar{b}''). \tag{21}$$

Since the pessimistic borrower undervalues honouring debts, his honour cut-off $\bar{y}(\cdot; \bar{b}')$ implicitly defined by

$$L(b', \bar{y}(b'; \bar{b}'); \bar{b}') = W_{\text{aut}}(\bar{y}(b'; \bar{b}')) \text{ for all } b' \tag{22}$$

provides an upper support function for the cut-off $y(\cdot)$ at $\bar{b}'$, depicted in Figure 8b. Since the pessimistic borrower's value function is differentiable, the implicit function theorem implies that $\bar{y}(\cdot; \bar{b}')$ is differentiable with $y_1(\bar{b}'; \bar{b}') > 1$ for all $\bar{b}' > b^*$.[16]

Thus far, we have established that the slope of the cut-off $y(\cdot)$ is zero approaching the risk-free limit b^* from the left, but greater than one approaching b^* from the right. Therefore, the cut-off has a downward kink at b^*, so it has no differentiable upper support function at this point. This means we have established:

Lemma 4. *At every $\bar{b}' \neq b^*$, there exists a differentiable upper support function $\bar{y}(\cdot; \bar{b}')$ for $y(\cdot)$, and hence a differentiable lower support function $\underline{q}(\cdot; \bar{b}')$ for $q(\cdot)$. Moreover, $y(\cdot)$ has an downward kink at b^* with $0 = y'(b^*-) < 1 < y'(b^*+)$.*

To construct a differentiable lower support function for W, we begin by constructing a differentiable lower support function for $W_{\text{hon}}(b', x')$. However, this time, we use a lazy borrower's value function that differs from the pessimistic value function used to construct (21). The lazy borrower correctly anticipates x', but incorrectly anticipates b' to be $\bar{b}'$. He takes on a debt of $\bar{b}''(x') = \hat{b}''(\bar{b}', x')$ independently of his previous obligation of b'. His value function is

$$M(b', x'; \bar{b}') = u(x' - b' + q(\bar{b}''(x'))\bar{b}''(x')) + \beta W(\bar{b}''(x')). \tag{23}$$

This means that,

$$\underline{W}(b'; \bar{b}') = W_{\text{aut}}(x') + \int_{\bar{y}(b'; \bar{b}')}^{x^{\max}} \left[M(b', x'; \bar{b}') - W_{\text{aut}}(x')\right] f(x')\, dx' \tag{24}$$

is a lower support function for W at $\bar{b}'$. We would like to establish that $\underline{W}(\cdot; \bar{b}')$ is differentiable. First, $M(\cdot, x'; \bar{b}')$ is continuously differentiable for all $(x', \bar{b}')$. Second, we note that without loss of generality, we may assume some optimal policy $\hat{b}''(\cdot, \cdot)$ is measurable, and

[16] Apply the implicit function theorem on the pessimistic borrower's value function to get

$$\bar{y}_1(\bar{b}'; \bar{b}') = \frac{u_1(\bar{c}'(\bar{b}', y(\bar{b}')))}{u_1(\bar{c}'(\bar{b}', y(\bar{b}'))) - u_1(x' - s)} > 1.$$

hence the resulting lazy policy $\bar{b}''(\cdot)$ is also measurable.[17] Third, the measurability of the lazy policy implies that $M_1(b', \cdot; \bar{b}')$ is measurable for all $(b', \bar{b}')$. Moreover, it is possible to show that $M_1(b', \cdot; \bar{b}')$ is uniformly bounded for all b' in some open neighbourhood of $\bar{b}'$. Hence the Leibniz rule for differentiating under the integral sign implies that $\underline{W}(\cdot; \bar{b}')$ is differentiable at $b' = \bar{b}'$ with[18]

$$\underline{W}_1(b'; \bar{b}') = \int_{\bar{y}(b';\bar{b}')}^{x^{\max}} M_1(\bar{b}', x'; \bar{b}') f(x')\, dx'. \tag{25}$$

This means we have established:

Lemma 5. *At every $\bar{b}'$, there exists a differentiable lower support function $\underline{W}(\cdot; \bar{b}')$ for W.*

First-Order Conditions Recipe. We can now return to the original problem (19). If $\hat{b}'$ is an optimal debt choice at the state (b, x), then it maximises

$$\phi(b'; b, x) = u(x - b + q(b')b') + \beta W(b'). \tag{26}$$

We now apply the steps of the recipe to derive first-order conditions. First, using $\underline{q}(\cdot; \bar{b}')$ and $\underline{W}(\cdot; \bar{b}')$, we can construct a differentiable lower support for this objective $\phi(\cdot; b, x)$ at any $(\bar{b}'; b, x)$. Second, the objective has a trivial upper support function. Third, the Differentiable Sandwich Lemma (Lemma 1) implies the borrower's objective is differentiable at the optimal debt choice $\hat{b}'$. Fourth, repeated application of the Reverse Calculus Lemma (Lemma 3) implies that y, q and W are differentiable at $\hat{b}'$.[19] We have thus proved:

Theorem 4. *Suppose $\hat{b}'(\cdot, \cdot)$ is an optimal policy function, fix any state (b, x), and set $\hat{b}' = \hat{b}'(b, x)$. If $\hat{b}' \neq b^*$, then the following first-order condition holds and the endogenous functions y, q and V that appear in it are differentiable at $\hat{b}'$:*

$$u_1(\hat{c}(b, x))(q(\hat{b}') + q_1(\hat{b}')\hat{b}') = \beta V_1(\hat{b}'), \tag{27}$$

[17] See the Measurable Maximum Theorem in Aliprantis and Border (2006, Theorem 18.19).

[18] See for example Weizsäcker (2008, Theorem 4.6).

[19] We repeatedly apply the Reverse Calculus Lemma (Lemma 3) as follows. First, we apply rule (i) (summation) to (26) to establish that both terms, $b' \mapsto u(x - b + q(b')b')$ and $b' \mapsto \beta W(b')$ are differentiable at $\hat{b}'$. Hence W is differentiable at $\hat{b}'$. Next, we apply rule (iv) (function composition) to $b' \mapsto u(x - b + q(b')b')$, which establishes that $b' \mapsto q(b')b'$ is differentiable at $\hat{b}'$. Next, we apply rule (ii) (multiplication) to establish that q is differentiable at $\hat{b}'$. This also means that both sides of (20b) are differentiable. So we apply rule (iv) (function composition) to the right side, $b' \mapsto \beta[1 - F(y(b'))]$ and conclude that y is differentiable at $\hat{b}'$.

where $\hat{c}(b,x) = x - b + q(\hat{b}'(b,x))\hat{b}'(b,x)$ *and*

$$V_1(\hat{b}') = \int_{y(\hat{b}')}^{x^{max}} u_1(\hat{c}(\hat{b}', x'))f(x')\,dx', \tag{28}$$

$$q_1(\hat{b}') = -\beta F_1(y(\hat{b}'))y_1(\hat{b}'), \tag{29}$$

$$y_1(\hat{b}') = \frac{u_1(\hat{c}(\hat{b}', y(\hat{b}')))}{u_1(\hat{c}(\hat{b}', y(\hat{b}'))) - u_1(y(\hat{b}') - s)}. \tag{30}$$

The first-order condition can be interpreted as follows. The borrower equates the marginal benefit of owing debt with the marginal cost. The marginal cost consists of the expected marginal utility of the foregone consumption when repaying the following period (when the endowment shock is above the default cut-off). The marginal benefit consists of the marginal utility of consumption times the marginal revenue from promising an extra payment to the lender. Reverse calculus allows us to quantify the marginal revenue of promises. Specifically, when promising an extra payment, the honour probability decreases according to q_1, which reflects an increase in the default cut-off of y_1.

Also note that the theorem gives formulae for all derivatives. In particular all prices and marginal prices can be written in terms of quantities. This means it is possible to write the first-order condition in terms of quantities only, which can be helpful in computational work.[20]

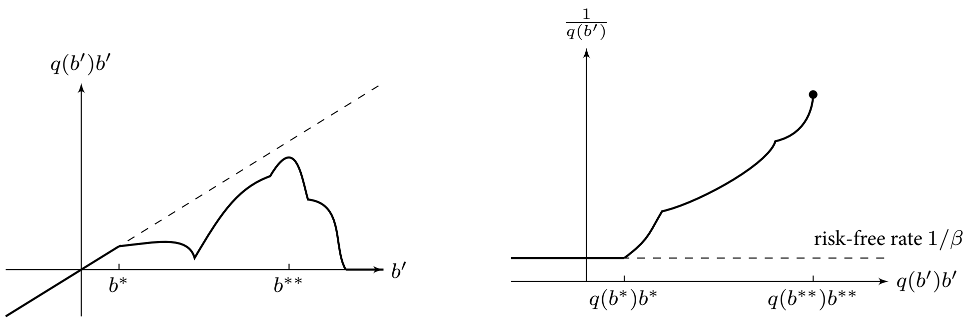

(a) Laffer curve for debt

(b) Endogenous interest rate

Figure 9: Characterisation of endogenous borrowing

Credit Limits. We now turn our attention to the borrower's behaviour near the credit limit. The amount the lender is willing to pay, $q(b')b'$ in return for a promise of b' is not an increasing function. This is because there are two types of empty promises: $b' = 0$, and b' so large it is never honoured. The borrower's return on promises therefore follows

[20] A FOC without prices is $u_1(\hat{c}(b,x))[1 - F(y(\hat{b}')) - F_1(y(\hat{b}'))y_1(\hat{b}')\hat{b}'] = V_1(\hat{b}')$.

a Laffer curve, depicted in Figure 9a. The borrower's credit limit is the maximum of this curve, $q(b^{**})b^{**}$, where

$$b^{**} = \arg\max_{b'} q(b')b'. \tag{31}$$

We apply the recipe for this new optimisation problem. If $b^{**} > b^*$, then we have already constructed a differentiable lower support function for q at b^{**}, so the Differentiable Sandwich Lemma (Lemma 1) together with the Reverse Calculus Lemma (Lemma 3) imply that q is differentiable at b^{**} with

$$q(b^{**}) + q_1(b^{**})b^{**} = 0. \tag{32}$$

Substituting this into the Euler equation (27), we see that the marginal benefit of taking on debt at b^{**} is zero, while the marginal cost is positive. Therefore, we conclude

Theorem 5. *For any given model primitives, either*

(i) the overall and risk-free credit limits coincide, i.e. $b^{**} = b^*$, *or*

(ii) the overall credit limit is higher and exhausting it is suboptimal, i.e. $b^{**} > b^*$ *and* $\hat{b}'(b, x) < b^{**}$ *for all states* (b, x).

This conclusion is a logical generalisation of behaviour in Aiyagari's (1994) model. Both here and there, the borrower reaches the risk-free credit limit with positive probability. In the model we study, the overall credit limit is potentially higher, as the borrower has the additional possibility of taking out risky loans. However, behaviour near the two credit limits is strikingly different. Below the risk-free limit, the interest rate $1/q(b')$ remains constant as the loan size $q(b')b'$ increases. Above the risk-free limit, the interest rate increases as the borrower takes on more debt and increases the default risk, as depicted in Figure 9b. This difference accounts for why borrowers might exhaust their risk-free limit, but not their overall limit.

Arellano (2008, Figure 2) plots a similar Laffer curve as in Figure 9a. Possibly for computational reasons, her curve is smooth and does not depict the upward kink of the Laffer curve at the risk-free limit, b^*. She does not apply first-order conditions along the Laffer curve.

Final Remarks. Our main contributions here are: (i) establishing that first-order conditions – involving derivatives of both policy functions and value functions – hold at optimal choices, (ii) providing formulas for these derivatives, so (marginal) prices can be written in terms of quantities, and (iii) applying the same logic to a Laffer curve for debt, concluding that borrowers do not exhaust their credit limits.

Some questions remain. First, we do not know if the Laffer curve is single-peaked. Second, the IID shock assumption was important for Arellano (2008) to establish that the default policy is a cut-off rule. More generally, persistent shocks cause interest rates to depend on the shock in addition to the size of the loan, which is crucial for understanding how credit markets operate when borrowers are distressed. Nevertheless, we believe our analysis can be generalised. Chatterjee, Corbae, Nakajima and Ríos-Rull (2007, Theorem 3) established that two-sided cut-off rules are optimal in an environment with persistent shocks. We conjecture that it is possible to construct differentiable support functions for the two cut-offs, and use this to construct a differentiable upper support function for the repayment probability.

5.2 Adjustment Costs

Firms are slow to adjust prices, labour forces, and capital stocks in reaction to changes in market conditions. One explanation for this is that firms face adjustment costs such as fixed costs or other non-convex costs. There is a large literature investigating how shocks propagate in the presence of adjustment costs and whether or not adjustment costs amplify shocks; see the surveys by Khan and Thomas (2008a), Leahy (2008), and Caplin and Leahy (2010). However, most of this literature is purely empirical, because the theory of adjustment costs faces two important obstacles. One is the complexity of optimal policy functions. Both theoretical and empirical analysis has only been tractable thus far when optimal polices involve smooth cut-off rules for determining when adjustments take place.[21] The other is the difficulty in deriving recursive first-order conditions, as the value of adjustment is not differentiable in general. Caballero and Engel (1999) use shocks that enter linearly into the production function to smooth out the kinks in the value function. Under this specific structure, they are able to take first-conditions to characterise optimal adjustments. To make this operational, they conjecture that adjustments follow a smooth two-sided (S, s) policy, but only verify this numerically.[22] Gertler and Leahy (2008) study a quadratic approximation of the firm's objective function in which the non-differentiable terms in the continuation value of adjustment vanish and optimal policies are smooth two-sided (S, s). They establish low error bounds for this approximation for an appropriate range of adjustment cost and shock parameters. Elsby and Michaels (2014) use first-order conditions under the conjecture that the optimal adjustment policy is a smooth two-sided (S, s) policy, also without providing sufficient conditions on primitives for this conjecture to hold. For the purposes of illustration, Cooper and Haltiwanger (2006, Section 3.2) and Khan and Thomas (2008b, Appendix B) provide derivatives only in the absence of fixed costs; we show these derivatives hold generally. An alternative approach is to assume that information arrives gradually over continuous time; see Harrison *et al.* (1983), Stokey (2008), and Golosov and Lucas (2007).

[21] Specifically, we say that a policy is a smooth two-sided (S, s) policy if (i) for every capital (or labour or price) level, the set of shocks for which the firm makes an adjustment is an interval and (ii) the upper and lower end points of this interval are differentiable functions of the capital level.

[22] Caballero and Engel (1999, Footnote 16)

The fundamental problem is that if a firm invests more today, then it might defer subsequent investment longer. Thus a small change in today's choice may lead to a lumpy change in a later choice, giving a non-differentiable and non-concave value of investment. We show that at optimal adjustment choices, the value function is differentiable so that recursive first-order conditions are applicable. We require only very weak assumptions on the primitives. In particular, our result remains true even when optimal policies are not two-sided (S, s) (see for example Bar-Ilan, 1990).

Model. In a general formulation, a firm is endowed with a capital stock k and shock z. Shocks evolve according to a Markov process with conditional distribution $P(z'|z)$. In each period, the firm's flow profit is $\pi(k,z)$; for example $\pi(k,z) = pf(k,z) - rk$ where p is output price, f is the production function, and r is the rental rate of capital. The firm pays an adjustment cost $c(k',k,z)$; non-adjustment is costless. We assume the flow profit $\pi(\cdot,z)$ is differentiable for all z, and that the adjustment cost $c(\cdot,\cdot,z)$ function is differentiable at all points (k,k',z) such that $k' \neq k$. For example, this accommodates the pure fixed-cost function, $c(k',k,z) = I(k' \neq k)$. The firm's value before adjusting its capital stock at state (k,z) is $V(k,z)$. Its value after adjusting its capital stock to k' is $W(k',z)$. These two value functions are related by the following two Bellman equations:

$$V(k,z) = \max_{k'} \pi(k,z) - c(k',k,z) + \beta W(k',z), \tag{33a}$$

$$W(k',z) = \int V(k',z')\, dP(z'|z). \tag{33b}$$

Our goal is to establish the first-order condition for the capital choice k'

$$c_1(k',k,z) = \beta W_1(k',z) \tag{34}$$

and to derive a formula for the marginal value of investment $W_1(k',z)$ at the optimal choice $k' = \hat{k}'(k,z)$. If there is a fixed cost of an adjustment, then this formula will only be satisfied when the agent makes an adjustment, i.e. at shocks z lying in the optimal adjustment set

$$\hat{A}(k) = \left\{ z : \hat{k}'(k,z) \neq k \right\}. \tag{35}$$

Ingredients: Differentiable Lower Support Functions. The most important ingredient of the recipe is providing a differentiable lower support function for the value function V. We consider a lazy manager who knows the optimal policy when he begins with a familiar capital stock of $k = \bar{k}$. The obvious lazy manager policy of sticking to the same capital choice when $k \neq \bar{k}$ is not useful here, because it leads to a discontinuous lazy value function.[23] Instead, we consider a lazy manager who uses the familiar adjustment set and

[23] This obvious lazy manager makes an extra adjustment even if the capital stock is only slightly different from the familiar level.

adjustment level for unfamiliar capital stocks, i.e. he waits until he draws a shock $z \in \hat{A}(\bar{k})$ and adjusts to $\hat{k}'(\bar{k}, z)$. Thereafter, his choices coincide with the rational manager. His value function is

$$L(k, z; \bar{k}) = \pi(k, z) + \begin{cases} \beta \int L(k, z'; \bar{k}) \, dP(z'|z) & \text{if } z \notin \hat{A}(\bar{k}), \\ -c(\hat{k}'(\bar{k}, z), k, z) + \beta W(\hat{k}'(\bar{k}, z), z) & \text{if } z \in \hat{A}(\bar{k}). \end{cases} \tag{36}$$

It is straightforward to calculate the lazy manager's marginal value of capital, because the capital stock k does not affect any subsequent choices:[24]

$$L_1(k, z; \bar{k}) = \pi_1(k, z) + \begin{cases} \beta \int L_1(k, z'; \bar{k}) \, dP(z'|z) & \text{if } z \notin \hat{A}(\bar{k}), \\ -c_2(\hat{k}'(\bar{k}, z), k, z) & \text{if } z \in \hat{A}(\bar{k}). \end{cases} \tag{37}$$

First-Order Conditions Recipe. If $\hat{k}'$ is an optimal choice at the state (k, z), then $\hat{k}'$ maximises

$$\phi(k'; k, z) = \pi(k, z) - c(k', k, z) + \beta W(k', z). \tag{38}$$

We follow the recipe to establish first-order conditions. First, substituting in (36) and (33b) into (38) gives a differentiable lower support function for $\phi(\cdot; k, z)$ at $\hat{k}'$. Second, Lemma 2 provides a differentiable upper support function. Third, Lemma 1 establishes the following theorem. (The reverse calculus step is trivial in this problem.)

Theorem 6. *If making an adjustment is optimal at state (k, z), i.e. $z \in \hat{A}(k)$, then the investment value W is differentiable in capital at $(\hat{k}'(k, z), z)$ and*

$$c_1(\hat{k}'(k, z), k, z) = \beta W_1(\hat{k}'(k, z), z) = \beta \int \tilde{L}_1(\hat{k}'(k, z'), z') \, dP(z'|z), \tag{39a}$$

$$\text{where } \tilde{L}_1(k, z) = \pi_1(k, z) + \begin{cases} \beta \int \tilde{L}_1(k, z') \, dP(z'|z) & \text{if } z \notin \hat{A}(k), \\ -c_2(\hat{k}'(k, z), k, z) & \text{if } z \in \hat{A}(k). \end{cases} \tag{39b}$$

The first equation says that the marginal adjustment cost should equal the marginal value of investment, which is the same for both rational and lazy managers. The second equation says that the marginal value of increasing investment equals the expected marginal increase in profit until the next adjustment plus the marginal decrease in the subsequent adjustment cost. We have thus shown that first-order conditions are generally valid even if the optimal adjustment policies are not (S,s). In other words, we have established that the applicability of first-order conditions is not an obstacle to the theoretical analysis of the implications of adjustment costs to prices, labour forces, and capital stocks. The only remaining obstacle is understanding when optimal policies are (S,s).

[24] The lazy manager's marginal value follows from the chain rule applied to (i) the expected discounted profit as a function of all state-contingent capital choices, holding adjustment times fixed, and (ii) the lazy capital choices as a function of initial capital k only.

5.3 Social Insurance

Governments run public health, unemployment and disability insurance programs, and private companies offer insurance contracts. These are constrained by frictions such as hidden information, adverse selection, and moral hazard. Informal insurance arises within well-connected families and communities when they can partially overcome these frictions. There is a large literature studying informal insurance, and the interaction of informal insurance with other forms of insurance.[25] In the dynamic insurance models of Thomas and Worrall (1988, 1990) and Kocherlakota (1996), the main issue is how cross-subsidisation may be self-enforcing. Agents with good luck subsidise those with bad luck in return for promises of future payments and insurance. These papers study smooth convex environments in which the Benveniste and Scheinkman (1979) theorem provides a formula for the marginal cost of making promises.[26] However, some important insurance problems involve non-smooth settings. We focus on a setting similar to that of Morten (2015), which is an extension of Ligon *et al.*'s (2002) model of self-enforcing dynamic insurance. Villagers share risk among themselves by both sharing divisible output and sending some members of the community to find temporary work in cities. The temporary migration decisions are inherently discrete as they involve a fixed cost of moving to the city and back. Other examples of indivisible items in village economies include livestock, medical treatments, agricultural land (due to high legal costs), and houses. This environment is non-smooth and non-concave, so the marginal cost of promises does not exist globally. Nevertheless, our envelope theorem applies and allows us to characterise optimal insurance policies in terms of the marginal cost of promises. Optimal policies involve sharing risk through allocating indivisible temporary work obligations; divisible consumption is then allocated to smooth out the marginal utility of consumption across states.

Model. Consider the following dynamic risk-sharing game between two households $h \in \{1, 2\}$. Each period begins with a Markov shock $s \in S$ with transition function $p(s'|s)$. The shock determines each household's endowment of a divisible consumption good, $C_h(s)$. The aggregate endowment is $C(s) = C_1(s) + C_2(s)$. In addition, each household may produce M units of the consumption good from temporary migrant work in a city. We write $d_h = 1$ if the household migrates, and $d_h = 0$ otherwise. We assume that the utility from consumption $u(\cdot, d_h)$ is differentiable, and that the marginal utility approaches infinity as consumption approaches zero. The autarky value of each

[25] Apart from the papers we discuss, Townsend (1994), Attanasio and Ríos-Rull (2000), and Krueger and Perri (2006) are important papers.

[26] Kocherlakota (1996) mistakenly claims his value function is differentiable. Koeppl (2006) amends his Bellman equation along the lines of Thomas and Worrall (1988). See also Ljungqvist and Sargent (2012, Chapter 20), and Rincón-Zapatero and Santos (2009, Section 4.2) for further discussion.

household is

$$V_h^{\text{aut}}(s) = \max_{d_h} u(C_h(s) + Md_h, d_h) + \beta \sum_{s'} p(s'|s) V_h^{\text{aut}}(s'). \tag{40}$$

Before investigating the social insurance arrangements with autarky constraints, we present the social planner's problem with Negishi weights η_1 and η_2:

$$W(s) = \max_{c_1, d_1} \eta_1 u(c_1, d_1) + \eta_2 u(c_2, d_2) + \beta \sum_{s' \in S} p(s'|s) W(s') \tag{41a}$$

$$\text{where} \quad c_1(s) + c_2(s) = C(S) + (d_1 + d_2)M. \tag{41b}$$

The first-order condition with respect to c_1 gives the Borch (1962) equation

$$\frac{u_1(c_1, d_1)}{u_1(c_2, d_2)} = \frac{\eta_2}{\eta_1}. \tag{42}$$

This means that after the social planner allocates the migration decisions, she adjusts the consumption good until the planner's marginal rate of substitution between the households is equal to the ratio of Negishi weights at all states and dates.

Now, we add in autarky constraints to study the optimal incentive-compatible social insurance contract. The value function for household 1 can be formulated recursively in terms of a principal-agent problem in which household 1 acts as an insurer and is able to promise future utility to household 2. This promised utility is a state variable, and has a corresponding promise-keeping constraint. Both households can leave the contract at any time, so there is an autarky constraint for each of them.

$$\begin{aligned}
V(s, v_2) = \max_{c_1, d_1, d_2, v_2'(s')} \; & u(c_1, d_1) + \beta \sum_{s' \in S} p(s'|s) V(s', v_2'(s')) \\
\text{s.t.} \quad (\text{PK}_2) \quad & u(c_2, d_2) + \beta \sum_{s' \in S} p(s'|s) v_2'(s') = v_2, \\
(\text{A}_1\text{-}s') \quad & V(s', v_2'(s')) \geq V_1^{\text{aut}}(s') \text{ for all } s' \in S, \\
(\text{A}_2\text{-}s') \quad & v_2'(s') \geq V_2^{\text{aut}}(s') \text{ for all } s' \in S, \\
\text{where} \quad & c_2 = C(s) + (d_1 + d_2)M - c_1.
\end{aligned} \tag{43}$$

First-Order Conditions. We need not apply the four steps of the recipe here, because the problem fits into the framework of Theorem 3. For the purposes of applying the theorem, the discrete choice is the migration allocation (d_1, d_2), and the continuous choices are the state-contingent promised utilities $v_2'(\cdot)$. The consumption choice c_1 can be eliminated by substituting in the promise-keeping constraint. Theorem 3 delivers the following:

Theorem 7. *Let $(\hat{d}_1(s, v_2), \hat{d}_2(s, v_2), \hat{v}_2'(s'|s, v_2))$ be an optimal choice at state (s, v_2). Fix some state (s, v_2) and some future shock s'. If neither autarky constraint binds for the choice of $\hat{v}_2'(s'|s, v_2)$, then the value function $V(s', \cdot)$ is differentiable at $\hat{v}_2'(s'|s, v_2)$ with*

$$-\frac{u_1(\hat{c}_1(s, v_2), \hat{d}_1(s, v_2))}{u_1(\hat{c}_2(s, v_2), \hat{d}_2(s, v_2))} = V_2(s', \hat{v}_2'(s'|s, v_2)) \tag{44}$$

$$= -\frac{u_1(\hat{c}_1(s', \hat{v}_2'(s'|s, v_2)), \hat{d}_1'(s', \hat{v}_2'(s'|s, v_2)))}{u_1(\hat{c}_2(s', \hat{v}_2'(s'|s, v_2)), \hat{d}_2'(s', \hat{v}_2'(s'|s, v_2)))}, \tag{45}$$

where $\hat{c}_1(s, v_2)$ and $\hat{c}_2(s, v_2)$ are implicitly defined in terms of the other optimal choices via the promising-keeping constraints.

This equation is the Borch (1962) equation which characterises perfect insurance – the social planner's marginal rate of substitution is equated across states and time periods. This means we have shown that with both divisible and indivisible choices, there is perfect insurance between households at all states and times for which the autarky constraints are lax. When an autarky constraint binds, the Negishi weights are adjusted and perfect insurance continues until an autarky constraint binds in the future. This generalises the conclusion drawn by Thomas and Worrall (1988) when indivisible choices are absent.

6 Conclusion

All envelope theorems have a sandwich idea at their core. Previous proofs were structured around sandwiches of inequalities of directional derivatives. By restructuring around sandwiches of differentiable upper and lower support functions, we gain two things. First, we do not require any of the strong technical conditions from previous envelope theorems, and can accommodate primitives with Inada conditions. Second and more importantly, our approach potentially applies to any type of endogenous functions that might need to be differentiated in a first-order condition.

There are potentially many ways to mix and match different constructions of upper and lower halves of sandwiches. We used five constructions throughout, namely (i) horizontal lines above maxima, (ii) supporting hyperplanes above concave functions, (iii) reverse calculus, (iv) lazy value functions below rational value functions, and (v) pessimistic cut-off rules. Of these constructions, only the reverse calculus construction is truly unprecedented. The power of our approach derives from the ability to combine these constructions, and the four-step recipe provides an intuitive way to organise them. For example, the unsecured credit application uses all but the supporting hyperplane construction. There are also other possibilities that we did not explore. Decision makers can be "lazy" in ways that lead to upper support functions, such as being lazily optimistic about future opportunities. In bargaining games, a lower support function for one player's value function leads to an upper support function for the other player's value function.

To conclude, our new approach reveals that trade-offs which previously seemed poorly behaved in fact have smooth structures within them that lead to first-order characterisations of optimal decisions.

A Further Reverse Calculus Rules

This appendix provides two further reverse calculus rules that were not used in the paper, but might be useful for other problems. Specifically, the rules relate to convex combinations and function composition.

The rule for convex combinations is complicated, because the forward calculus step is not obvious. The following lemma incorporates both a forward and reverse calculus result.

Lemma 6. *Suppose $E : C \to [0,1]$, $F : C \to \mathbb{R}$, and $G : C \to \mathbb{R}$ have differentiable lower support functions e, f, and g respectively at $\bar{c}$. Consider the function*

$$H(c) = E(c)F(c) + (1 - E(c))G(c).$$

If $F(\bar{c}) > G(\bar{c})$, then

(i) The function $h(c) = e(c)f(c) + (1 - e(c))g(c)$ is a differentiable (local) lower support function for H at $\bar{c}$.

(ii) If H is differentiable at $\bar{c}$, then e, f, and g are also differentiable at $\bar{c}$.

Proof. Consider the two functions,

$$\begin{aligned} h(c) &= e(c)f(c) + (1 - e(c))g(c) \\ \tilde{h}(c) &= E(c)f(c) + (1 - E(c))g(c). \end{aligned}$$

Since $f(\bar{c}) > g(\bar{c})$, we have that $h(c) \leq \tilde{h}(c)$, and hence $h(c) \leq H(c)$ in some open neighbourhood of $\bar{c}$. This establishes part (i).

For part (ii), we see that $\tilde{h}$ is differentiably sandwiched between h and H at $\bar{c}$. By the Differentiable Sandwich Lemma, $\tilde{h}$ is differentiable at $\bar{c}$. This implies $E(c) = [\tilde{h}(c) - g(c)]/[f(c) - g(c)]$ is also differentiable at $\bar{c}$. Therefore, both terms of H, namely $E(c)F(c)$ and $(1 - E(c))G(c)$, have differentiable lower support functions, $E(c)f(c)$ and $(1 - E(c))g(c)$, respectively. Part (i) of Lemma 3 implies that both terms are differentiable at $\bar{c}$, and hence F and G are differentiable at $\bar{c}$. ☐

Finally, we consider function composition of two endogenous functions.

Lemma 7. *If $H(c) = J(K(c))$ is differentiable at $\bar{c}$, where*

- *$J : \mathbb{R} \to \mathbb{R}$ has an inverse J^{-1} and a differentiable lower support function $j(\cdot)$ at $K(\bar{c})$,*
- *$K : \mathbb{R} \to \mathbb{R}$ has an inverse K^{-1} and a differentiable lower support function $k(\cdot)$ at $\bar{c}$, and*
- *$j'(K(\bar{c})) \neq 0$ and $k'(\bar{c}) \neq 0$,*

then J and K are differentiable at $K(\bar{c})$ and $\bar{c}$ respectively.

Proof. We assume without loss of generality that $j'(K(\bar{c})) > 0$.[27] We now establish that this implies j^{-1} is a differentiable *upper* support function for J^{-1}. To see this, we evaluate the inequality $j(c) \leq J(c)$ at $J^{-1}(x)$ which gives

$$j(J^{-1}(x)) \leq J(J^{-1}(x)) = x.$$

Applying j^{-1} to both sides gives $J^{-1}(x) \leq j^{-1}(x)$.

We can express $K(\cdot)$ as a function of J and H as follows:

$$J^{-1}(H(c)) = J^{-1}(J(K(c))) = K(c).$$

This has a differentiable upper support function $j^{-1}(H(c))$ at $\bar{c}$. Thus K has differentiable upper and lower support functions at $\bar{c}$, and is therefore differentiable by Lemma 1. Next, evaluating $H(c) = J(K(c))$ at $c = K^{-1}(x)$ gives

$$H(K^{-1}(x)) = J(K(K^{-1}(x))) = J(x),$$

so J is differentiable at $K(\bar{c})$ by the chain rule and inverse function theorem. □

References

Aguiar, M. and Gopinath, G. (2006). Defaultable debt, interest rates and the current account. *Journal of International Economics*, **69** (1), 64–83.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, **109** (3), 659–684.

Aliprantis, C. D. and Border, K. C. (2006). *Infinite Dimensional Analysis*. Springer Verlag, 3rd edn.

Amir, R., Mirman, L. J. and Perkins, W. R. (1991). One-sector nonclassical optimal growth: Optimality conditions and comparative dynamics. *International Economic Review*, **32** (3), 625–644.

[27] If $j'(K(\bar{c})) < 0$, then the lemma can be applied to $\tilde{H}(c) = \tilde{J}(K(c))$, where $\tilde{H}(c) = -\frac{1}{H(c)+x_0}$ and $\tilde{J}(c) = -\frac{1}{J(c)+x_0}$, where x_0 is a suitable constant to prevent division by zero near $K(\bar{c})$. In this case, $\tilde{j}(c) = -\frac{1}{j(c)+x_0}$ is a lower support function for $\tilde{J}$ with a strictly positive derivative $\tilde{j}'(c) = \frac{1}{[j(c)+x_0]^2} j'(c)$.

Araújo, A. and Scheinkman, J. A. (1977). Smoothness, comparative dynamics, and the turnpike property. *Econometrica*, **45** (3), 601–620.

Arellano, C. (2008). Default risk and income fluctuations in emerging economies. *American Economic Review*, **98** (3), 690–712.

— and Ramanarayanan, A. (2012). Default and the maturity structure in sovereign bonds. *Journal of Political Economy*, **120** (2), 187–232.

Attanasio, O. and Ríos-Rull, J.-V. (2000). Consumption smoothing in island economies: Can public insurance reduce welfare? *European Economic Review*, **44** (7), 1225–1258.

Bar-Ilan, A. (1990). Trigger-target rules need not be optimal with fixed adjustment costs: A simple comment on optimal money holding under uncertainty. *International Economic Review*, **31** (1), 229–234.

Benveniste, L. M. and Scheinkman, J. A. (1979). On the differentiability of the value function in dynamic models of economics. *Econometrica*, **47** (3), 727–732.

Borch, K. (1962). Equilibrium in a reinsurance market. *Econometrica*, **30** (3), 424–444.

Bulow, J. and Rogoff, K. (1989). Sovereign debt: Is to forgive to forget? *American Economic Review*, **79** (1), 43–50.

Caballero, R. J. and Engel, E. M. R. A. (1999). Explaining investment dynamics in U.S. manufacturing: A generalized (S,s) approach. *Econometrica*, **67** (4), 783–826.

Caplin, A. and Leahy, J. (2010). Economic theory and the world of practice: A celebration of the (S, s) model. *Journal of Economic Perspectives*, **24** (1), 183–201.

Chatterjee, S., Corbae, D., Nakajima, M. and Ríos-Rull, J.-V. (2007). A quantitative theory of unsecured consumer credit with risk of default. *Econometrica*, **75** (6), 1525–1589.

Clarke, F. H. (1975). Generalized gradients and applications. *Transactions of The American Mathematical Society*, **205**, 247–262.

Clausen, A. and Strub, C. (2016). Introduction to economic theory, lecture notes.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment cost. *Review of Economic Studies*, **73**, 611–633.

Eaton, J. and Gersovitz, M. (1981). Debt with potential repudiation: Theoretical and empirical analysis. *Review of Economic Studies*, **48** (2), 289–309.

Elsby, M. and Michaels, R. (2014). Fixed adjustment costs and aggregate fluctuations, mimeo.

Gertler, M. and Leahy, J. (2008). A Phillips curve with an Ss foundation. *Journal of Political Economy*, **116** (3), 533–572.

Golosov, M. and Lucas, R. E. (2007). Menu costs and Phillips curves. *Journal of Political Economy*, **115** (2), 171–199.

Harrison, J. M., Sellke, T. M. and Taylor, A. J. (1983). Impulse control of brownian motion. *Mathematics of Operations Research*, **8**, 454–466.

Hatchondo, J. C. and Martinez, L. (2009). Long-duration bonds and sovereign defaults. *Journal of International Economics*, **79** (1), 117–125.

Khan, A. and Thomas, J. K. (2008a). Adjustment costs. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*, Basingstoke: Palgrave Macmillan.

— and — (2008b). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, **76** (2), 395–436.

Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *Review of Economic Studies*, **63** (4), 595–609.

KOEPPL, T. V. (2006). Differentiability of the efficient frontier when commitment to risk sharing is limited. *B.E. Journal of Macroeconomics*, **6** (1), 1–6.

KRUEGER, D. and PERRI, F. (2006). Does income inequality lead to consumption inequality: Evidence and theory. *Review of Economic Studies*, **76**, 163–193.

KRUGER, A. Y. (2003). On Fréchet subdifferentials. *Journal of Mathematical Sciences*, **116**, 3325–3358.

KYDLAND, F. E. and PRESCOTT, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, **85** (3), 473–491.

LEAHY, J. (2008). s-S models. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*, Basingstoke: Palgrave Macmillan.

LIGON, E., THOMAS, J. P. and WORRALL, T. (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of Economic Studies*, **69** (1), 209–244.

LJUNGQVIST, L. and SARGENT, T. J. (2012). *Recursive Macroeconomic Theory*. MIT Press, 3rd edn.

MILGROM, P. and SEGAL, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, **70** (2), 583–601.

MIRMAN, L. J. and ZILCHA, I. (1975). On optimal growth under uncertainty. *Journal of Economic Theory*, **11** (3), 329–339.

MORTEN, M. (2015). Temporary migration and endogenous risk sharing in village India, mimeo.

RINCÓN-ZAPATERO, J. P. and SANTOS, M. S. (2009). Differentiability of the value function without interiority assumptions. *Journal of Economic Theory*, **144** (5), 1948–1964.

ROCKAFELLAR, R. T. (1970). *Convex Analysis*. Princeton University Press.

— and WETS, R. J. (1998). *Variational Analysis*. Springer.

SANTOS, M. S. (1991). Smoothness of the policy function in discrete time economic models. *Econometrica*, **59** (5), 1365–1382.

STOKEY, N. L. (2008). *The Economics of Inaction: Stochastic Control models with fixed costs*. Princeton University Press.

— and LUCAS, R. E. (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

THOMAS, J. and WORRALL, T. (1988). Self-enforcing wage contracts. *Review of Economic Studies*, **55** (4), 541–553.

— and — (1990). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory*, **51** (2), 367–390.

TOWNSEND, R. M. (1994). Risk and insurance in village India. *Econometrica*, **62** (3), 539–591.

VARIAN, H. R. (1992). *Microeconomic Analysis*. Norton, 3rd edn.

WEIZSÄCKER, H. v. (2008). Basic measure theory, lecture notes.