Direct Complementarity

Jonathan Weinstein*

Washington University in St. Louis

First version: April 2017; This version: July 2020

Abstract

I propose a novel notion of complementarity in consumer theory, called *direct complementarity*, for which I provide four equivalent definitions. I point out a novel critique of the leading definition of complementarity, which is based on cross-price effects: such effects are sensitive to changes in the basis used to describe the space of available bundles. Direct complementarity, on the other hand, is defined for preferences over an abstract vector space of bundles, without reference to a particular basis (or list of "goods"), and provides a consistent definition across all pairs of composite goods, i.e. linear combinations of standard goods. Direct complementarity better captures the intuitive notion of one good's effect on the value of another, hence the term *direct*; cross-price effects are best understood as an *indirect* consequence of such relationships.

*I thank attendees at the 2017 North American Summer Meetings of the Econometric Society and the 2017 Games and Economic Behavior conference in honor of Ehud Kalai, the Fall 2018 NBER conference on Mathematical Economics, and the spring 2020 conference on Complementarity at ICERM as well as John Nachbar and Brian Rogers, for useful feedback. This version is preliminary, and in particular needs more citations; the main sources for key concepts in the standard theory of demand are the graduate microeconomics textbooks of Kreps and of Mas-Colell, Whinston, Green.

1 Introduction

The notions of complementarity and substitutability are central to our understanding of preference and choice in multidimensional environments. The simplest formal context for these notions is a preference over a product $X \times Y$ of ordered sets, represented by a utility function u. In this context, we ordinarily say that quantities x and y are strict complements if the function u is strictly *supermodular*: that is, if for every pair $x_1 \geq x_2 \in X$, $u(x_1, y) - u(x_2, y)$ is strictly increasing in y. If the domain is $\mathbb{R}^2$ and u is twice-differentiable, then strict supermodularity of u is equivalent to positivity of the cross-partial derivative: $u_{xy} > 0$. That is, the marginal value of x is increasing in y, and vice-versa. This, of course, has important consequences for optimization. Under the reverse of this condition, we say that x and y are substitutes.

In the earliest days of consumer theory, a positive cross-partial derivative was indeed used to define complementarity of goods (ALPE). However, the definition was found as early as 1915 (Stigler & Boulding 1915) to be problematic, for the fundamental reason that consumer preferences are customarily taken to be *ordinal*, with a multiplicity of interchangeable cardinal representations, and the sign of second derivatives is not preserved by monotone transformations. As customary, assume that the consumer's preference $\succeq$ on $\mathbb{R}^n$ (the space of bundles of n goods) is represented by a twice-differentiable function $u : \mathbb{R}^n \to \mathbb{R}$. Denote partial derivatives by u_i, u_{ij}, etc. Assume, as usual, that $u_i(x) > 0$ for all x and i, i.e. all coordinates represent goods. If $f : \mathbb{R} \to \mathbb{R}$ is any strictly increasing smooth function, then $v = f \circ u$ is another equally valid representation. Then, applying the chain rule twice, we obtain for any i, j:

$$v_{ij} = f' u_{ij} + f'' u_i u_j \tag{1}$$

so that clearly if f'' is large enough and positive, all goods will appear to be complements under representation v, while the opposite will be true if f'' is extremely negative.

Accordingly, it has long been customary in consumer theory to define complements and substitutes by looking at demand functions rather than at utility. If the compensated cross-price effect is negative – that is, if a price increase in good i, along with a change in income which keeps utility constant, decreases the demand for good j – we say that goods i and j are *Hicksian complements*. It is shown in standard textbooks that (unlike gross complementarity, which looks at uncompensated cross-price effects) this definition is symmetric.

Samuelson (1974) gives an excellent survey of the history of complementarity. He complains at length that the Hicksian definition does not relate directly to natural intuitions about complementarity, but matches them at best indirectly. He explains clearly the problems caused by non-cardinal utility, but argues that nonetheless, complementarity between, say, sugar and tea "ought" to mean something like the marginal utility from additional sugar and tea being greater than the sum of the individual marginal utilities. The direction of cross-price effects, on the other hand, reflects the intuitive idea of complementarity in a roundabout way at best. I call this "Samuelson's complaint." He then proposes defining complementarity by fixing a distinguished utility representation, "money-metric" utility, and then using the straightforward criterion based on the cross-partial of this utility. He calls this "money-metric complementarity." This proposal, though, has gained little attention.[1]

In this paper, I hope to make two important contributions to the under-

[1] His paper has over 600 citations; I scanned many of the citing papers, and almost all were referencing the survey component of his work and not his novel definition. Also interestingly, I presented versions of this paper to a variety of audiences over several years, during which time I was completely unaware that my primary definition is similar to Samuelson's. No one pointed out the similarity, or brought up Samuelson's work at all. One could view my paper as resurrecting Samuelson's forgotten idea, providing an array of equivalent characterizations which validate its robustness, and providing a novel formal "complaint" about cross-price effects to complement his intuitive complaint.

standing of complementarity in consumer theory. One concerns the popular Hicksian definition. To complement Samuelson's intuitive complaints about this definition, I make the following, more formal complaint: the definition is not invariant to changes in coordinates on the space of bundles, and should not even be properly viewed as a well-defined relation on goods. My meaning here will be elucidated shortly, with the example and discussion in Section 2.

With this critique in mind, the next contribution is to give a collection of equivalent definitions for a concept I call *direct complementarity*. One way to define direct complementarity is to distinguish a representation, which must be chosen locally, and look at its cross-partials, similarly to Samuelson. In the case of quasi-linear preferences, this representation is the standard one, and the definition is equivalent to Samuelson's. Perhaps more compellingly, direct complementarity can be uniquely characterized by a few simple principles, or axioms. One meta-principle is that complementarity should be defined between any two vectors in bundle-space, with the usual complementarity between two goods corresponding to the case of basis vectors. Once we expand our horizons in this way, it is natural to require that complementarity should not depend on the coordinate system chosen for bundle-space; as mentioned, this criterion is failed by Hicksian complementarity, but it is satisfied by direct complementarity. This "basis-free" principle is enforced implicitly, by simply giving definitions which refer to an abstract vector space of bundles without a distinguished basis. The explicit axioms which pin down the definition can be described informally as follows:

1. **Symmetry**: Complementarity (and substitutability) should be symmetric relations.

2. **Linearity**: If v is a complement of both w and z, it also complements all linear combinations of w and z.

3. **Unanimity**: If for vectors v and w, the cross-partial of utility in di-

rections v, w has the same sign for *all* smooth utility representations, we should classify v and w as complements or substitutes accordingly.

4. **Neutrality**: If moving in the v^* direction does not alter any marginal rates of substitution, then v^* is neutral (neither complement nor substitute) to all vectors.

The usefulness of Axiom 3(and a reason it may have been overlooked) is revealed by inspecting equation (1). As mentioned earlier, *if* i and j are *goods*, with positive marginal utility, (1) tells us that their cross-partial will vary arbitrarily as we change representations. However, if at least one of them is a *neutral*, with zero marginal utility, their cross-partial is consistent across representations. When we expand complementarity to a relation on any pair of vectors, this is a powerful observation. Since the neutrals have codimension 1, if we are given some non-neutral v^*, any vector is a linear combination of v^* and a neutral. This may give the reader a glimmer of why, under the proper regularity conditions, the axioms pin down a unique definition of complementarity. My main hope for this paper is that the reader comes to share my view that the collection of equivalent definitions for direct complementarity mark it as *the* natural concept.

ROADMAP

2 A key example

The following example will illustrate the basis dependence of Hicksian complementarity and the contrast with direct complementarity.

Alice is forming a portfolio from one safe asset (asset 0) and three risky assets (1, 2 and 3.) There are no short-sale constraints. Her preferences are given by

$$u(x_0, x_1, x_2, x_3) = x_0 + x_1 + x_2 + x_3 - \frac{x_1^2}{2} - \frac{x_2^2}{2} - \frac{x_3^2}{2}$$

For motivation, note that preferences of this form will arise if the risky assets are identically, independently and normally distributed, and Alice has CARA preferences. Then, the certainty equivalent of her portfolio will be of the same form as u, a quadratic with no cross-terms. Of course, in the reduced form above this utility could also apply to consumer goods.

At prices $(p_0, ..., p_3)$, demand for each risky asset is given by $x_i^* = 1 - p_i/p_0$, with remaining wealth invested in the safe asset. The three risky assets hence have no cross-price effects. Notice also that because utility is quasi-linear in Asset 0, income effects apply only to Asset 0, so that we do not need to know Alice's wealth, or worry about compensated vs. uncompensated cross-price effects, when evaluating Hicksian complementarity of Assets 1-3. For convenience we henceforth fix $p_0 = 1$, i.e. there is a risk-free interest rate of 0.

Now suppose Asset 3 is replaced by a mutual fund, M, containing one-third of a share of each of assets 1, 2, and 3, so Alice will be forming her portfolio from assets 0, 1, 2, and M. Prices $q_0 = 1, q_1, q_2, q_M$ will be such that Alice faces the very same optimization problem; the change of coordinates is given by

$$q_1 = p_1, q_2 = p_2, q_M = \frac{p_1 + p_2 + p_3}{3}$$

The quantities z_i Alice must purchase in order to form an equivalent portfolio are given by the linear transformation

$$z_0 = x_0, z_1 = x_1 - x_3, z_2 = x_2 - x_3, z_M = 3x_3$$

Rewriting Alice's demand in the new coordinate system, i.e. using z and q, gives

$$z_1 = -2q_1 - q_2 + 3q_3, z_2 = -q_1 - 2q_2 + 3q_3, z_M = 3q_1 + 3q_2 - 9q_3 + 3$$

So, in the new coordinate system, Assets 1 and 2 have a negative cross-price effect and are Hicksian complements, though they are defined identically as before, the full set of available bundles is unchanged, and the price of any bundle is unchanged. Why does this happen? In the new coordinate system, the meaning of two key concepts has subtly changed. An increase in the price of Asset 1 or 2 has changed meaning, and so has an increase in demand for Asset 1 or 2. When we talk about a price increase, it is always implicit that the price of other goods is held fixed. The change in the definition of Asset 3, to Asset M, changes the meaning of a price increase in Asset 1, because a different *ceteris* is *paribus*. When we increase the price of Asset 1 with Assets 2 and M fixed, the meaning in the original coordinates is that p_1 increases, p_2 is fixed, and p_3 decreases. Still in terms of the original coordinates, this causes a decrease in x_1 and increase in x_3, with x_2 held fixed. So why do we see a cross-price effect on Asset 2? Because though Asset 2 is the same asset, $z_2 \neq x_2$, and holding x_2 fixed is not the same as holding z_2 fixed. To increase x_3 with x_2 fixed, in the z-coordinates we must increase z_M while decreasing z_2.

To help gain further insight, let's rewrite Alice's utility in terms of the z_i:

$$v(z_0, z_1, z_2, z_3) = z_0 + z_1 + z_2 + z_M - \frac{z_1^2}{2} - \frac{z_2^2}{2} - \frac{z_M^2}{6} - \frac{z_1 z_M}{3} - \frac{z_2 z_M}{3}$$

In this formulation, it is intuitive that Asset 1 and Asset 2 are each substitutes for Asset M, since the corresponding cross-partials v_{1M}, v_{2M} are negative – and of course from a risk-management point of view, this makes perfect sense, because Asset M is positively correlated with each of Assets 1 and 2. The cross-price effect, in the $z-q$ coordinate system, between Assets 1 and 2 should be thought of as "indirect." An increase in the price q_1 of Asset 1 causes an increase in demand for its substitute, Asset M, which is in turn a substitute for Asset 2, so there is a decrease in demand for Asset 2. On the other hand, $v_{12} = 0$, just as $u_{12} = 0$, corresponding to the independent returns of Assets 1 and 2, so under a definition based on this cross-partial, Assets 1 and 2 would be neither complements nor substitutes, in contrast with their status as Hicksian complements. In fact, the cross-partial $u_{12} = 0$ is preserved by any redefinition of assets other than 1 and 2.

To sum up: Hicksian complementarity is sensitive to the entire basis used to span bundle space, and not just the two goods being considered and the larger space of available bundles. This reinforces Samuelson's complaint that Hicksian complementarity doesn't match our intuitions about how two goods may enhance each other's consumption; in fact, if we allow the coordinate system to vary, it is not really a well-defined relation on the two goods at all! In this paper, an important desideratum for complementarity is that it should depend only on the two goods under consideration and on preferences over bundle space, not on the coordinate system for bundle space – we call this "basis-free". Now, once a utility function is fixed, its cross-partials automatically satisfy the condition of being basis-free, but of course we then must reckon with the problem that the choice of utility function matters. Ultimately, one of our definitions for the general case will give, for any preference, a distinguished utility function, chosen locally around each point. Our first definition will focus on the case of quasi-linear preferencds.

2.1 Direct Complementarity: The Quasi-Linear Case

In the special case of quasi-linear preferences, our new concept can be defined very simply, in terms of the standard utility representation:

Definition 1. *Let preferences be quasi-linear, represented by $u(x) = x_0 + f(x_1, \ldots, x_k)$. Then, at a bundle x, we say that two goods $i, j \neq 0$ are* direct complements *if $u_{ij}(x) > 0$ and* direct substitutes *if $u_{ij}(x) < 0$.*

Note that any other representation of the same preferences which takes *this functional form* would differ from u only by a constant, so there is no ambiguity in the definition.

To further explicate the distinction between direct complements and Hicksian complements, recall the characterization of compensated cross-price effects as cross-partials of the expenditure function:

$$\frac{\partial h_i(p,u)}{\partial p_j} = \frac{\partial^2 E(p,u)}{\partial p_i \partial p_j}$$

where h_i is Hicksian demand and E is the expenditure function (the minimum cost of achieving utility u at prices p.) In this formulation, it is clear that Hicksian complementarity is a binary relation on *price* vectors, *not* on goods – more specifically, on the infinitesimal changes in price represented by ∂p_i and ∂p_j. Crucially, price vectors p do *not* lie in the space of bundles of goods; rather, the set of price vectors is the *dual* space to the set of bundles, i.e. the set of linear functions from bundles to $\mathbb{R}$. When, as customary, we speak of the relation of Hicksian complementarity as if it is a relation on goods, we are implicitly making use of an isomorphism between bundle-space and its dual, price-space. As should be familiar from linear algebra, this isomorphism is non-canonical, i.e. basis-dependent. The change ∂p_i in price-space is defined by increasing the price of i with the price of other goods *in the basis* held fixed. In the example, when Asset 3 is replaced by Asset M, changing *one* basis element in bundle-space, *all* basis vectors in price-space change meanings: As discussed earlier, ∂q_1 represents an increase in the price of Asset 1 with the prices of Assets 2 and M held fixed, while ∂p_1 of course assumed the prices of Assets 2 and 3 were held fixed. On the other hand, the value of u_{12} depends only on the definitions of Assets 1 and 2.

The vital point here is that, fundamentally, direct complementarity is a symmetric relation between *goods*, while Hicksian complementarity is a symmetric relation between *price changes*. Standard textbook language, by treating Hicksian complementarity as a relation between *goods*, confounds a vector space with its dual.

Later, we will define direct complements and substitutes for general preferences. Our Definition A is a natural extension of the quasi-linear case. We observe that near a given point, *any* smooth utility function looks quasi-linear (up to a second-order approximation) under an appropriate change of coordinates, where the direction of the income effect takes the role of Good 0, and we extend the definition accordingly. Other equivalent definitions, which don't refer to the quasi-linear case or to any representation, will help make it seem that this concept is inevitable.

3 Definitions and Preliminary Results

One often gains clarity by working in greater generality and abstraction. In this paper, rather than think about complementarity only for pairs of goods from $\{1, 2, ..., n\}$, we will define it simultaneously for all pairs of vectors (v, w) in bundle-space, with the special case of basis vectors (e_i, e_j) corresponding to the usual complementarity between goods i and j. Though not strictly necessary, we actually gain greater elegance by extending the concept of complementarity to any (single) element of the tensor product $V \otimes V$, with complementarity of v and w corresponding to the complementarity of $v \otimes w$. It will turn out that the set of complementary tensors is a half-space in $V \otimes V$. Also, instead of working with one representation for a preference, it is natural to work with the set of all functions representing the same preference, as we will below.

Let V be a real vector space of finite dimension n; we write V and not $\mathbb{R}^n$ so as not to designate a favored basis. As with existing definitions, we will always treat complementarity as a local property around a given point $x \in V$. The point x will be fixed throughout this section, and we will sometimes omit x from notation when it is convenient and causes no confusion. We use the standard notation V^* for the dual space, i.e. the set of linear functions from V to $\mathbb{R}$. When $u : V \to \mathbb{R}$ is differentiable at x, we write $Du(x) \in V^*$ for the "Fréchet derivative" of u at x, meaning that $Du(x)(v)$ is the first-order approximation around x of $u(x + v) - u(x)$. (If we fixed a basis of V and wrote $Du(x)$ in terms of the dual basis of V^*, it would be the usual gradient of u at x.) We will need preferences to be well-behaved near x:

Assumption 1 (Smooth Representation)**.** *Preferences are represented by a function u which is C^∞ on a neighborhood around the point $x \in V$ of interest, and satisfies $Du(x) \neq 0$.*

All representations u referenced hereafter will be assumed to satisfy Assumption 1. The two halves of the assumption could be called smoothness and non-degeneracy. Sometimes, for emphasis, I call such a u a *smooth* representation; this refers to both parts of the assumption. Note that it is always easy to represent the same preferences with a $\hat{u}$ which violates the assumptions, by letting $\hat{u} = f \circ u$ for an increasing $f : \mathbb{R} \to \mathbb{R}$ which either has $f'(u(x)) = 0$, or f non-differentiable at $u(x)$.

Write $u \sim \hat{u}$ if $\hat{u} = f \circ u$ for a C^∞ function $f : \mathbb{R} \to \mathbb{R}$ with $f' > 0$ everywhere.

Definition 2. *For any C^∞ utility function $u : V \to \mathbb{R}$, the equivalence class $[u] = \{\hat{u} : \hat{u} \sim u\}$ is an* ordinal C^∞ function*.*

The derivative of $[u]$, $D[u]$, is defined as the set of derivatives $\{D\hat{u} : \hat{u} \in [u]\}$. The chain rule says that $D\hat{u}(x) = f'(u(x))Du(x)$, so

$$D[u](x) = \{\alpha Du(x) : \alpha \in (0, \infty)\}$$

i.e. the derivative is defined up to positive scalar. Then $D[u](x)$ corresponds, canonically, to a half-space in V, namely the directions corresponding to increased utility, or "goods," $G(x) = \{v \in V : Du(x)(v) > 0\}$, bounded by the "indifference plane," $I(x) = \text{Ker}(Du(x))$. By standard linear algebra, all functionals which are positive on the half-space G differ only by a positive scalar, hence are in $D[u](x)$. So G carries all the first-order information available from the ordinal function $[u]$. Note that, given a choice of coordinates, the hyperplane I tells us all marginal rates of substitution (MRS) between any two goods i and j; the MRS is the λ such that $e_i - \lambda e_j \in I$.

At each $x \in V$, $D^2u(x)$ can he viewed as a (symmetric) bilinear form $D^2u(x) : V \times V \to \mathbb{R}$, where in a coordinate system, $D^2u(x)(e_i, e_j)$ would be the usual cross-partial in coordinates (i, j). An equivalent formulation we will find useful is to write $D^2u(x) \in (V \otimes V)^*$, i.e. a linear functional on 2-tensors over V. Again, let $\hat{u} = f \circ u$, then

$$D^2\hat{u}(x)(v, w) = f'(u(x))D^2u(x)(v, w) + f''(u(x))Du(x)(v)Du(x)(w) \quad (2)$$

$$D^2\hat{u}(x) = f'(u(x))D^2u(x) + f''(u(x))[Du(x) \otimes Du(x)]$$

where the second line follows by definition of the form $[Du(x) \otimes Du(x)]$. Then we can write

$$D^2[u] = \{\alpha D^2u + \beta(Du \otimes Du) : \alpha \in (0, \infty), \beta \in \mathbb{R}\}$$

The second term in the last equation, parametrized by β which can have either sign, is what challenges our ability to define complementarity, by making the sign of $D^2u(v, w)$ vary across the equivalence class $[u]$. This motivates us to define the "first-order-indifferent tensors," on which this term vanishes:

$$I^2 := \text{Ker}(Du \otimes Du) \subseteq (V \otimes V)$$

Note that I^2 is a space of codimension one (i.e, dimension $n^2 - 1$) in $(V \otimes V)$. An alternate characterization is:

$$I^2 = \text{Span}(I \otimes V \cup V \otimes I)$$

where the inclusion $\supseteq$ is immediate, and the reverse holds by counting dimensions. That is, I^2 is the span of tensors $v \otimes w$ where either v or w is a first-order neutral. When restricted to I^2, $D^2[u]$ is well-defined up to a positive scalar. Conversely, any form which, when restricted to I^2, is a positive scalar multiple of D^2u, is an element of $D^2[u]$, i.e.

$$D^2[u] = \{\nu \in (V \otimes V)^* : \exists \alpha \in \mathbb{R}^+ : \forall t \in I^2 : \nu(t) = \alpha D^2u(t)\}$$

The reverse inclusion holds because of the standard fact that $Du(x) \otimes Du(x)$ spans the forms which annihilate its kernel, I^2.

So $D^2[u]$ determines, and is determined by, a half-space in I^2 which we call the "complementary tensors," $C_I(x) := \{t \in I^2 : D^2u(x)(t) > 0\}$, bounded by the "neutral tensors," $N_I(x) := I^2 \cap \mathrm{Ker}(D^2u)$, and we call the other half of I^2 the "substitutive tensors" $S_I = -C_I$. Note that $D^2[u]$ provides no information in its behavior outside of I^2 – such behavior is constrained only by linearity and is otherwise completely representation-dependent. In particular, if we are working with a basis where every coordinate vector is a "good," i.e. $e_i \in G$, then, by (2), for large positive f'' we have $D^2\hat{u}(e_i, e_j) > 0$ for all i, j, and for very negative f'' the reverse inequality. That is, as asserted in the introduction, $[u]$ contains both supermodular and submodular functions. The intuition here is that we can "see" second-order effects only when first-order effects are zero. The lack of cardinal utility confounds our ability to distinguish the second-order phenomenon of complementarity, except when first-order effects vanish.

To summarize, around a point x, the first-order information about an ordinal function $[u]$ can be characterized by the half-space of goods, G, bounded by the indifference plane I. The first derivative also determines the first-order-indifferent tensors I^2. Then, the second-order effects partition I^2 into C_I, N_I and S_I. To make the interpretation of this partition more clear, let $v_1 \in I$, $v_2 \in G$, so $v_1 \otimes v_2 \in I^2$. Then, what does it mean to have $v_1 \otimes v_2 \in C_I$? It means that a small movement in the v_2 direction changes v_1 from a neutral to a good, i.e. $Du(x + \varepsilon v_2)(v_1) > 0$ for small enough positive ε. To be even more concrete, let e_1, e_2, e_3 be basis vectors which are all in G, and let

$$v_{12} := \frac{e_1}{Du(x)(e_1)} - \frac{e_2}{Du(x)(e_2)} \in I$$

be the vector, unique up to a scalar, which is in both the indifference plane and the plane spanned by e_1 and e_2. Then $v_{12} \otimes e_3 \in C_I$ means that motion in the e_3 direction makes v_{12} a good – which is the same as saying that it increases the relative value of Good 1 to Good 2, i.e. the marginal rate of substitution between the goods. We interpret this to mean that Good 3 complements Good 1 *better* than it complements Good 2. This *relative* complementarity – equivalent to complementarity between Good 3 and the "neutral" v_{12} – is robust to changes in representation.

The following geometric interpretation is illuminating: when preferences are locally strictly convex around x, $D^2u(x)(v, v) < 0$ for all $v \in I$, i.e. D^2u is negative-definite on I, so $-D^2u$ is an inner product on I, and is unique up to scalar as we let u vary over $[u]$. A pair $(v_1, v_2) \in I$ are complements in our sense if $-D^2u(v_1, v_2) < 0$, i.e. if the angle between them according to this inner product is greater than $\pi/2$, neutrals if they are orthogonal, and substitutes if the angle is less than $\pi/2$. In particular, any good is a substitute for itself, as is intuitive under convexity.

3.1 The "Money" Good

As in the last section, we will be focusing on complementarity around a fixed point $x \in V$. To recap, we found that when at least one of w, z is a neutral, $D^2u(w, z)$ has the same sign for all smooth representations u, and we can use this sign to define complementarity. The remaining question is how best to extend this definition to any pair of goods. The key will be to designate one vector, not in the indifference plane, which is neutral (not a complement or substitute) with respect to all other goods. In the quasi-linear case of Section 2.1, this was Good 0 (money). What can fill in for money in the general case? Such a vector must be neutral with respect to all goods in the indifference plane. Fortunately, such a vector always exists; generically there will be only one choice with this property, up to scalar. In the standard terminology of demand theory, this vector, which we call v^*_x, is simply the direction of the income effect in bundle space, i.e. the direction with no first-order effect on any marginal rate of substitution, as shown in Proposition 1. Note that as usual, the marginal rate of substitution between two goods $w, z \in V$ is defined by

$$MRS_{w,z} = \frac{Du(w)}{Du(z)}$$

for any smooth representation u.

Proposition 1. *The following are equivalent properties of a vector* $v_x^* \in V$*:*

1. *For each* $w \in I$*,* $v_x^* \otimes w \in N_I$*, i.e. for every smooth representation* u*,* $D^2u(v_x^*, w) = 0$*.*
2. v^* *satisfies* $D(MRS_{w,z})(v^*) = 0$ *for all* w, z *with* $z \notin I$*.*

Proof. In (2), note that $Du(z) \neq 0$ by assumption. Using the quotient rule and cancelling the denominator, $D(MRS_{w,z})(v^*) = 0$ is equivalent to

$$Du(z)D^2u(v_x^*, w) - Du(w)D^2u(v_x^*, z) = 0$$

By bilinearity of D^2u, this is equivalent to:

$$D^2u(v_x^*, Du(z)w - Du(w)z) = 0 \tag{3}$$

and by linearity of Du, $Du(Du(z)w - Du(w)z) = 0$, i.e. $Du(z)w - Du(w)z \in I$. Therefore, 1 implies that equation (3) holds for every w, z and hence implies 2.

To prove $1 \Rightarrow 2$, start with any $w \in I$, and apply equation (3) for any $z \notin I$ (such z exists by Assumption 1). The equation reduces to $D^2u(v_x^*, Du(z)w) = 0$, and since $Du(z) \neq 0$ this gives the desired result.

■

There is a dimension-counting argument that a non-zero v_x^* with these properties must exist: For any smooth u, D^2u induces a map from V to I^*, defined by $v \mapsto (w \mapsto D^2u(v, w))$. Since V has dimension n while I^* has dimension $n - 1$, the map has non-trivial kernel, and any element of the kernel has the desired property. It turns out to be inconvenient if the space of such elements happens to be more than one-dimensional, or happens to lie in I, so we introduce the following regularity assumption, which holds generically in the space of possible first and second derivatives of u:

Assumption 2. *The element $v_x^* \in V - \{0\}$ satisfying the conditions in Proposition 1 is unique up to scalar, and furthermore $v_x^* \notin I$.*

This v_x^* plays a key role in all of our coming definitions.

4 Direct Complementarity

In this section, we define what it means for two vectors $v, w \in V$ to be *direct complements*. To strengthen the intuition for this definition, we provide four equivalent definitions. The preliminary results in the previous section provide a foundation for understansing these definitions and why they are equivalent.

4.1 Definition A: Via a locally quasilinear representation

SECTION NEEDS REVISION Without referring to representation, the defining property of quasi-linear preferences is that there is a good ("money") with no impact on any marginal rates of substitution (MRS) between any two goods. Since we define complements locally, it makes sense to think about whether this property holds locally, i.e. there is a good with no first-order effect on any MRS. As per our discussion in the last section, such a vector always exists. From this "local quasilinearity" of preferences, we can construct a "locally quasilinear" representation, as defined here:

Definition 3. *A representation u^q is locally quasilinear at x if there is a $v_x^* \in V - \{0\}$ such that $D^2 u^q(v_x^*, v) = 0$ for all $v \in V$.*

Lemma 1. *For any preferences with a smooth representation u in a neighborhood around x, there is a locally quasilinear representation u^q at x.*

Proof. Once we identify v_x^* as described in Proposition 1, construct u^q by letting $u^q = f \circ u$ for any f satisfying:

$$\frac{f''(u(x))}{f'(u(x))} = -\frac{D^2u(x)(v_x^*, v_x^*)}{(Du(x)(v_x^*))^2}$$

To prove the claimed property of u^q, first note that it holds for $v = v_x^*$ by application of (2). Also, for $v \in I$, it holds by definition of v_x^*. Finally, v_x^* and I together span V (by the last clause in the definition of v_x^*), and the property extends linearly. ■

Then, by analogy with the quasilinear case discussed earlier, we say:

Definition A. Let u^q be a locally quasilinear representation at x. Then w, z are *direct complements* at x if

$$D^2u^q(x)(w, z) > 0$$

They are neutrals if this derivative is 0, and substitutes if it is negative.

Please note that there is apparent ambiguity in this definition, due to the multiplicity of locally quasilinear representations, but it is only apparent. Theorem 1 will show Definition A to be equivalent to other definitions regardless of the choice among such representations, so that the choice does not matter. Note that while we did not use Assumption 2 directly in this section, it is needed for the coming alternate definitions, and is needed for the claim that Definition A is not sensitive to the choice of u^q.

4.2 Definition B: Decomposition into Nutrients and Flavor

The next definition is created as follows: Note that Assumption 1 implies that I has dimension $n-1$, and Assumption 2 states that $v_x^* \notin I$. Therefore, any bundle w can be decomposed, uniquely, in the form

$$w = \lambda_w v_x^* + w^n$$

where $w^n \in I$ is first-order neutral. Conceptually, bundles are composed of **nutrients** (utility-rich at first-order, second-order-neutral) and *flavor* (first-order-neutral, with second-order impact). We now say

Definition B. w, z are *direct complements* at x if, when w and z are decomposed as above, $D^2u(x)(w^n, z^n) > 0$ for *all* smooth representations u, or, equivalently, for *any* smooth representation u. They are neutrals if this derivative is 0, and substitutes if it is negative.

The equivalence between "all" and "any" in the definition follows from our earlier observation that the sign of $D^2u(x)(w, z)$ is invariant to representation when w or z is first-order neutral.

Conceptually, it is the first-order-neutral "flavors" (w^n, z^n) which may be complements or substitutes, while v_x^* is simply a flavorless gruel of utility, complementing any bundle equally well. Direct complementarity of (w, z) is equivalent to complementarity of the "flavors," (w^n, z^n), a condition which is representation-invariant, as mentioned earlier. We also observe that the definition would remain equivalent if we replaced (w^n, z^n) by (w^n, z) or (w, z^n); this is immediate from bilinearity of D^2u and the defining property of v_x^*.

4.3 Definition C: An Axiomatic Approach

The notion of direct complements and substitutes can also be characterized by a list of natural axioms, which were summarized intuitively in the introduction:

Theorem 1. *Given a smooth preference defined on a neighborhood of a point $x \in V$, there is a unique partition of V^2 into three sets C, S, N satisfying the following axioms:*

1. ***Symmetry**: For all $(v, w) \in V^2$, (v, w) and (w, v) are in the same category.*

2. ***Linearity**: If $(v, w) \in C$ and $(v, z) \in C \cup N$, then $(v, \alpha w + \beta z) \in C$ for any $\alpha > 0$,$\beta \geq 0$. The same condition holds if S or N replaces C throughout.*

3. ***Unanimity**: If*

$$D^2 u(v, w) > 0$$

for all *smooth representations u, then $(v, w) \in C$. Similarly, if the cross-partial is 0 for all u then $(v, w) \in N$, and if negative then $(v, w) \in S$.*

4. ***Neutrality**: If v^* satisfies*

$$D(MRS_{v,w})(v^*) = 0$$

for all v, w with $w \notin I$, then $(v^*, v) \in N$ for all $v \in V$.

where all derivatives are taken at the point x.

The proof of Theorem 1 is in Appendix A. The idea is that the unanimity axiom pins down the definition when at least one of (v, w) is a first-order neutral, the neutrality axiom pins it down when one equals v^*, and linearity and symmetry do the rest, because, as noted, neutrals and v^* span V. Note that rather than making symmetry a separate axiom, we could have written symmetric versions of linearity and neutrality, and the same result would hold. I felt that the axioms were a little easier to read this way. Note that the adjective "smooth" in Axiom 3 is crucial; without it, the axiom is vacuous, because there is always a representation, violating Assumption 1, with $Du = D^2u = 0$.

With this theorem in hand, we naturally make the following definition:

Definition C. w, z are *direct complements* at x if $(w, z) \in C$ where C is as described in Theorem 1. Similarly, they are substitutes if $(w, z) \in S$ and neutrals if $(w, z) \in N$.

4.4 Definition D: Via Marginal Rates of Substitution

The vector v_x^*, which would be a consumer's marginal consumption under an increase in income, is in some sense a natural numeraire with which to judge the value of other goods. That is, we could consider the value of good w to be MRS_{w,v_x^*}. Accordingly, when we ask whether good z complements good w, it is natural to ask whether moving in direction z increases the value of good w in this sense. This leads to the following definition:

Definition D. We say w, z are *direct complements* at x if

$$D(MRS_{w,v_x^*})(z) > 0$$

Notice while on its face, this definition does not treat w and z symmetrically, we will prove that it is equivalent to the other definitions, hence *de facto* symmetric. Note also that this symmetry would *not* hold if we used an arbitrary numeraire in place of v_x^*.

4.5 Equivalence Theorem

The main result of the paper is:

Theorem 2. *Definitions A, B, C, and D are equivalent.*

The full proof is in the appendix. The basic ideas are as follows: Equivalence between Definitions A and B is almost immediate, because for the representation u^q used in A, $D^2u^q(x)(w, z) = D^2u^q(x)(w^n, z^n)$. To show that B is equivalent to C, we only need to verify that Definition B satisfies the axioms. Equivalence between A and D falls out almost immediately from the definition of MRS via u^q and the defining property of v_x^*.

5 Relationship between Direct and Hicksian Complementarity

Here, we'll elucidate the relationship between direct and Hicksian complementarity. This is simplest to discuss in the quasi-linear case; keep in mind that our discussion in Section 4.1 tells us that any case can, locally, be transformed into the quasi-linear by a change of basis (which only requires replacing one good with a suitable numeraire.) We'll fix here a basis $e_0, ..., e_k$ for V (letting $k = n - 1$) and a quasi-linear utility function $u(x) = x_0 + f(x_1, \ldots, x_k)$. Let W be the span of $e_1, ..., e_k$. We can view $D^2f(x)$ as a bilinear form on W, or equivalently as a function from W to W^*. Represent $D^2f(x)$ as a $k \times k$ matrix H, using the dual basis for W^*. This is nothing but the ordinary Hessian matrix. A little thought shows that cross-price effects are given by the matrix H^{-1}, so

- Goods i, j are direct complements iff $H_{ij} > 0$.
- Goods i, j are Hicksian complements iff $H_{ij}^{-1} < 0$.

This helps us understand the example from the introduction. The Hessian in the original coordinate system is just $H = -I$ where I is the identity, so that both H and H^{-1} are diagonal, and no goods are complements or substitutes in either sense. To translate to the new coordinates, we define a change-of-coordinate matrix

$$C = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$$

and write

$$\hat{H} = CHC^T = \begin{pmatrix} -1 & 0 & -\frac{1}{3} \\ 0 & -1 & -\frac{1}{3} \\ -\frac{1}{3} & -\frac{1}{3} & -\frac{1}{3} \end{pmatrix}$$

The change in coordinates affects only the third row and column of H, leaving $\hat{H}_{12} = H_{12} = 0$ unchanged; this would be true for any H, exemplifying the fact that the change in coordinates does not affect direct complementarity of Goods 1 and 2. It *does*, though, affect H^{-1}_{12}, which becomes $\hat{H}^{-1}_{12} = -1$, so that as calculated earlier, Goods 1 and 2 are now Hicksian complements. Please note that by setting $H_{12} = -\varepsilon$, we would easily obtain an example where Goods 1 and 2 change from (Hicksian and direct) substitutes to Hicksian complements (but still direct substitutes) under the change of coordinates.

References

Samuelson, Paul. 1974. "Complementarity: An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory." *Journal of Economic Literature* 12:1255–1289.

Stigler, G.J. & K.E. Boulding. 1915. On the Theory of the Budget of Consumer. In *Readings in Price Theory*, ed. O. Moeschlin & D. Pallaschke. Chicago: Richard D. Irwin.

A Proof of Theorem 1

Existence follows from the fact that the partition into C, S, N given by Definition A satisfies the axioms. In fact, the first three axioms would hold if any representation were used. Symmetry and linearity come from the fact that second derivatives are always symmetric bilinear forms, and unanimity is trivial. The neutrality axiom is clear from the definition of local quasilinearity, along with Assumption 2.

As for uniqueness: Let U be the set of pairs which are characterized identically for all partitions satisfying the axioms; we must show that $U = V^2$. From our preliminary results and the unanimity axiom, we know that $(v, w) \in U$ if $v \in I$ or $w \in I$. From the neutrality and symmetry axioms, we know that $(v^*, v) \in U$ and $(v, v^*) \in U$, where v^* is the unique element described in Assumption 2. Also, recall that by this assumption, $I \cup \{v^*\}$ spans V. So, given arbitrary (v, w), write

$$w = \lambda_w v_x^* + w^n$$

B Proof of Theorem 2

$A \Leftrightarrow B$: Let the representation u^q be as in Section 4.1, and given any $w, z \in V$, let each be decomposed as in Section 4.2. Then

$$\begin{aligned} D^2u^q(x)(w, z) &= D^2u^q(x)(\lambda_w v_x^* + w^n, \lambda_z v_x^* + z^n) \\ &= D^2u^q(x)(w^n, z^n) \end{aligned}$$

by bilinearity of D^2u^q and the defining property of u^q, and the result follows. Please note that as promised, this implies that Definition A cannot be sensitive to the choice of u^q among locally quasi-linear representations.

$A \Leftrightarrow C$: Having proved Theorem 1, it suffices to show that Definition A satisfies the four axioms. The first three axioms would actually be satisfied even if an arbitrary fixed smooth representation replaced u^q in Definition A. Symmetry and Linearity follow from symmetry and bilinearity of D^2u^q, and Unanimity holds trivially. Neutrality does require local quasilinearity of u^q.