Resource Consciousness in Classical Logic

Andreas Blass

Using Herbrand's Theorem, we define simple Herbrand validity, a sort of resource consciousness that makes sense in classical predicate logic. We characterize the propositional formulas all of whose first-order instances are simply Herbrand valid. The characterization turns out to coincide with a known characterization of game semantical validity for multiplicative formulas.

1 Introduction

Resource consciousness is one of the main features of linear and affine logic (3). In these logics, validity of a sequent $\Gamma \vdash A$ means that the conclusion A can be obtained from the list Γ of hypotheses using each hypothesis exactly once (in the case of linear logic) or at most once (in the case of affine logic). Thus, in these logics, a hypothesis that is used twice must occur twice in the list Γ. Classical logic is not resource conscious in this sense, because the tautology $A \to A \wedge A$ allows us to produce as many copies as we want from one hypothesis A.

Nevertheless, we shall show in this paper that a form of resource consciousness can be found in classical first-order logic. Using it, we shall define two strong notions of validity, one for first-order formulas and one for propositional formulas. The new notion of validity in first-order logic is not very well behaved and plays only an auxiliary role, leading up to the propositional notion. The propositional notion, on the other hand, is well behaved, admits a syntactic characterization, and coincides with game semantical validity, as defined in (2), for the multiplicative fragment of affine logic.

*Address: Mathematics Department, University of Michigan, Ann Arbor, MI 48109-1109, U.S.A. ablass@umich.edu
Partially supported by a grant from Microsoft Research

Submitted for
Proceedings of LLC9
Ingrid van Loon, Grigori Mints, Reinhard Muskens (eds.)

Except for this introduction and a concluding section, we shall not need any concepts from linear logic or game semantics; the body of this paper is entirely about classical logic.

2 An Example

Consider the sentence τ expressing the transitivity of a binary relation $<$,

$$(\forall x, y, z)\,(x < y \wedge y < z \longrightarrow x < z),$$

and the sentence λ expressing a longer version of transitivity,

$$(\forall x, y, z, w)\,(x < y \wedge y < z \wedge z < w \longrightarrow x < w).$$

Of course, τ logically implies λ. It is intutively clear that any proof of λ from τ will have to use the hypothesis τ twice, because any particular instance of λ (specific values of x, y, z, w) depends on two instances of τ (either the x, y, z and x, z, w instances or else the y, z, w and x, y, w instances). We shall make this intuition precise and use it to introduce a resource conscious version of logical consequence in which each hypothesis is to be used at most once.

The discussion of τ and λ made essential use of the fact that both of them are universal sentences, so that there is a clear notion of "instance" and we can count how many instances of the hypotheses are needed to establish one instance of the conclusion. One can, of course, give a "dual" discussion in the case of existential sentences. Here the issue will not be the number of uses of the hypothesis but the number of "guesses" of witnesses for the conclusion. Consider, for example, the contrapositive of $\tau \longrightarrow \lambda$, with the negations moved into the scopes of the quantifiers. From

$$(\exists x, y, z, w)\,(x < y \wedge y < z \wedge z < w \wedge \neg x < w)$$

follows

$$(\exists x, y, z)\,(x < y \wedge y < z \wedge \neg x < z).$$

Given a witness x, y, z, w for the hypothesis, we know that either x, y, z or x, z, w serves as a witness for the conclusion, but we do not know which. Our resource conscious version of logical consequence will prohibit this sort of "two attempts at the conclusion" as well as its dual, "two uses of the hypothesis."

The ideas described in this transitivity example can be easily carried over to similar situations involving universal sentences and existential sentences, but it is not immediately clear how to extend them to sentences involving quantifier alternations. The tool that makes such an extension possible is Herbrand's theorem, which we review, in a convenient form, in the next section.

3 Herbrand's Theorem

Herbrand's theorem associates to each first-order sentence ϕ a quantifier-free formula ϕ_H such that ϕ is valid if and only some disjunction of instances of ϕ_H is a tautology.

Convention 3.1 We work in first-order logic without equality. The propositional connectives are $\wedge$, $\vee$ and $\neg$. Negation is applied only to atomic formulas; if we write $\neg$ in any other context, it is to be understood as an "abbreviation" for the result of pushing the negation in to the atomic level, using De Morgan's laws and the usual interchange rules for quantifiers and negations. If we write $A \longrightarrow B$, it is to be understood as an abbreviation of $\neg A \vee B$. We also assume that every vocabulary under consideration has at least one constant symbol.

Remark 3.2 It would make no difference if we used the other common definition of $A \longrightarrow B$ as $\neg(A \wedge \neg B)$, for pushing the negation in to the atomic level would convert this to $\neg A \vee B$. On the other hand, the two common ways of eliminating the biconditional $A \longleftrightarrow B$, namely $(A \longrightarrow B) \wedge (B \longrightarrow A)$ and $(A \wedge B) \vee (\neg A \wedge \neg B)$ are genuinely different. Since we shall be dealing with notions of validity more restrictive than ordinary logical validity, it is not correct to identify logically equivalent formulas, as is often done. See Remark 5.3 below for more about the difference between the two versions of the biconditional. Except for that remark, we shall not use biconditionals again in this paper, so we need not address the question of which interpretation is to be preferred. We thank the referee for pointing out the need for caution on this point.

Definition 3.3 The *Herbrand form* ϕ_H of a first-order sentence ϕ is obtained by the following five steps.

1. Rename bound variables in ϕ so that no variable is quantified twice.

2. Pull out all the universal quantifiers into a prefix of second-order quantifiers, using the usual prenex operations plus the following equivalences to pull universal quantifiers past existential ones.

$$\begin{aligned} \exists y\, \forall x\, \alpha(x,y) &\quad\longleftrightarrow\quad \forall X\, \exists y\, \alpha(X(y),y) \\ \exists y\, \forall X\, \alpha(X(\vec{z}),y) &\quad\longleftrightarrow\quad \forall X\, \exists y\, \alpha(X(y,\vec{z}),y). \end{aligned}$$

3. Pull out all the existential quantifiers from the first-order part of the formula, using the usual prenex operations. At this point the formula consists of a block of universal second-order quantifiers (over functions), then a block of existential first-order quantifiers, and then a quantifier-free matrix.

4. Delete the universal quantifiers. The function symbols that they quantified are to be added to the vocabulary. So the formula has become an existential first-order sentence in the enlarged vocabulary.

5. Delete the existential quantifiers. The variables $\vec{x}$ that they quantified are the free variables of the Herbrand form ϕ_H, so we often write $\phi_H(\vec{x})$.

Notice that, if a universal quantifier in ϕ lies in the scopes of k existential quantifiers, then the corresponding function symbol in ϕ_H is k-ary, and its arguments, wherever it appears in ϕ_H, are the variables of those k existential quantifiers. If $k = 0$, then we have a 0-ary function symbol, i.e., a constant symbol, in ϕ_H.

Example 3.4 If ϕ is

$$\forall x\, \exists y\, \forall z\, P(x,y,z) \longrightarrow \forall x\, \exists y\, \forall z\, Q(x,y,z),$$

then, remembering that implication is treated as an abbreviation and negations are to be pushed in to the atomic level, we find that the universally quantified variables are the y in the antecedent and the x and z in the consequent. Thus, the Herbrand form ϕ_H is (if we use u, v, w as the new variables in the consequent at step 1)

$$P(x, Y(x), z) \longrightarrow Q(U, v, W(v)).$$

In stating Herbrand's Theorem, we use the usual terminology: "Tautology" means valid in propositional logic. The statement of the theorem also involves the notion of a "closed instance" of a formula. Since three other notions of "instance" will play a role later in the paper, we define all four notions here to avoid confusion.

Definition 3.5

- A *closed instance* of a first-order formula ϕ is obtained by substituting closed terms for all the free variables in ϕ.

- A *first-order instance* of a propositional formula ϕ is obtained by replacing the propositional variables in ϕ by first-order sentences.

- A *propositional instance* of a propositional formula ϕ is obtained by replacing the propositional variables in ϕ by propositional formulas.

- A *variable-merging instance* of a propositional formula ϕ is obtained by replacing the propositional variables in ϕ by propositional variables, not necessarily distinct.

If, as a result of the replacing involved in this definition, the negation symbol is applied to a non-atomic formula, then, in accordance with Convention 3.1, it is understood that the negations are to be pushed in until they apply to atomic formulas.

Theorem 3.6 (Herbrand) *A first-order sentence ϕ is valid if and only if some finite disjunction of closed instances of ϕ_H is a tautology.*

Proof The first three steps in the construction of ϕ_H produce formulas logically equivalent to ϕ. (In the case of the rules for pulling universal quantifiers past existential ones, the soundness of the rules may be easier to see in the dual form,

$$\forall y\,\exists x\,\beta(x,y) \longleftrightarrow \exists X\,\forall y\,\beta(X(y),y),$$

a version of the axiom of choice.) At step 4, where we obtain the sentence $\exists\vec{x}\,\phi_H$, the equivalence is lost, since the truth value of this sentence may depend on the interpretations of the function symbols that had previously been universally quantified. Nevertheless, it is clear that step 4 does not affect logical validity; thus ϕ is logically valid if and only if $\exists\vec{x}\,\phi_H$ is.

Like any existential sentence (in a vocabulary with at least one constant symbol), $\exists\vec{x}\,\phi_H$ is valid if and only if it is true in every "term model," i.e., in every structure where every element is the value of a closed term. The reason is that every structure has a substructure that is a term model, and truth of existential sentences is preserved upward from substructures to superstructures.

A term model amounts to an assignment of truth values to all atomic sentences. So $\exists\vec{x}\,\phi_H$ is true in every term model if and only if every such truth assignment makes at least one closed instance of ϕ_H true. Finally, by the compactness theorem for propositional logic, this is equivalent to the existence of finitely many closed instances of ϕ_H such that every truth assignment makes at least one of them true, i.e., such that their disjunction is a tautology. $\square$

Remark 3.7 Like Herbrand (4) but unlike some modern presentations such as (7), we defined the Herbrand form of a sentence without first putting it into (first-order) prenex form. Prenex operations would, in general, move some universal quantifiers into the scopes of more existential quantifiers and would thus complicate the Herbrand form. In the proof of our main theorem, it will be important that these complications do not occur.

Unlike both (4) and (7) we have presented Herbrand's theorem as a semantical result, characterizing validity. Herbrand did not work with semantics, and he presented his theorem as a characterization of provability in a certain deductive system. Of course the completeness theorem implies that the two points of view are equivalent, but Herbrand's proof of his theorem was nearly simultaneous with Gödel's proof of the completeness theorem.

Example 3.8 In the previous example,

$$\forall x\, \exists y\, \forall z\, P(x,y,z) \longrightarrow \forall x\, \exists y\, \forall z\, Q(x,y,z)$$

is not logically valid, but it becomes so if we write a second P in place of Q. No disjunction of closed instances of the Herbrand form $P(x,Y(x),z) \longrightarrow Q(U,v,W(v))$ is a tautology (as one can make all atomic formulas with P true and all those with Q false). But if we replace Q with P, the new Herbrand form $P(x,Y(x),z) \longrightarrow P(U,v,W(v))$ has a single instance that is a tautology; replace the variables x, z, and v by the closed terms U, $W(Y(U))$, and $Y(U)$, respectively.

Example 3.9 In the transitivity example from Section 2, the Herbrand form of $\tau \longrightarrow \lambda$ is

$$(x < y \wedge y < z \longrightarrow x < z) \longrightarrow (T < U \wedge U < V \wedge V < W \longrightarrow T < W).$$

Here x, y, z are variables and T, U, V, W are constant symbols. No single closed instance of this formula is a tautology, but there are two closed instances whose disjunction is a tautology. For example, replace x, y, z by T, U, V for the first instance and by T, V, W for the second. (Alternatively, replace x, y, z by U, V, W for the first instance and by T, U, W for the second.)

Similarly, for the (rephrased) contrapositive

$$(\exists x, y, z, w)\,(x < y \wedge y < z \wedge z < w \wedge \neg x < w) \longrightarrow \\ (\exists x, y, z)\,(x < y \wedge y < z \wedge \neg x < z),$$

the Herbrand form is

$$(T < U \wedge U < V \wedge V < W \wedge \neg T < W) \longrightarrow (x < y \wedge y < z \wedge \neg x < z),$$

and we obtain a tautologous disjunction of two closed instances by the same substitutions as before.

Notice that the closed instances used in this example correspond exactly to the two uses of the hypothesis and the two attempts at the conclusion in the proofs of $\tau \longrightarrow \lambda$ and its contrapositive, as discussed in Section 2.

4 Simple Herbrand Validity

The preceding example suggests measuring resource usage, or at least bounding it from below, by the number of closed instances needed to produce a tautologous Herbrand disjunction. That is, if, as in the example, two disjuncts are needed, then we regard this as indicating that a hypothesis was used twice or that two attempts were made to obtain witnesses for a conclusion. The following definition is, therefore, intended to capture the idea that no such duplication is needed.

Definition 4.1 A first-order sentence ϕ is *simply Herbrand valid* if some closed instance of its Herbrand form ϕ_H is a tautology.

Thus, for example, $\forall x\, \exists y\, \forall z\, P(x,y,z) \longrightarrow \forall x\, \exists y\, \forall z\, P(x,y,z)$ is simply Herbrand valid, but $\tau \longrightarrow \lambda$ is not (where τ and λ are as in Section 2). The reader is invited to check that $\tau \wedge \tau \longrightarrow \lambda$ is simply Herbrand valid. Repeating the hypothesis τ makes available, in a single instance of the Herbrand form, the two instances of τ that we need.

Although simple Herbrand validity clearly has something to do with using hypotheses only once (and, dually, making only one guess at a witness for a conclusion), it depends too heavily on the quantifier structure to serve as a really good model of resource consciousness. It is conscious of (and prohibits) the sort of reuse of hypotheses that manifests itself in the occurrence of different instances, but it ignores other sorts of reuse. Thus, for example, $\forall x\, (P(x) \longrightarrow P(x) \wedge P(x))$ is simply Herbrand valid yet expresses the antithesis of resource consciousness. (A simpler example, if there are no objections to 0-ary predicate symbols, is $P \longrightarrow P \wedge P$.)

The situation improves if we abstract from the quantifier structure in the following way. Instead of considering a particular first-order formula, consider all first-order instances of some propositional formula.

Definition 4.2 A propositional formula is *universally simply Herbrand valid*, abbreviated *usHv*, if all its first-order instances are simply Herbrand valid.

The following example indicates that this notion captures the idea of resource consciousness better than simple Herbrand validity does. Contrast it with the observation above that $A \longrightarrow A \wedge A$ is simply Herbrand valid.

Example 4.3 $A \longrightarrow A \wedge A$ is not usHv. Indeed, its first-order instance

$$\forall x \, \exists y \, P(x,y) \longrightarrow (\forall x \, \exists y \, P(x,y) \wedge \forall x \, \exists y \, P(x,y))$$

is not simply Herbrand valid. Its Herbrand form is

$$P(x, Y(x)) \longrightarrow P(U,v) \wedge P(W,z).$$

No single closed instance of this is a tautology. Indeed, in any closed instance, the antecedent matches at most one conjunct from the consequent, so there is a truth assignment verifying the antecedent and falsifying a different conjunct in the consequent. (There are two closed instances whose disjunction is a tautology.)

5 Universal Simple Herbrand Validity

The purpose of this section is to prove the main result of the paper, a syntactic characterization of universal simple Herbrand validity.

Definition 5.1 A propositional formula is *binary* if no propositional variable occurs in it more than twice.

Theorem 5.2 *For any propositional formula ϕ, the following are equivalent.*

1. *ϕ is universally simply Herbrand valid.*
2. *ϕ is a propositional instance of a binary tautology.*
3. *ϕ is a variable-merging instance of a binary tautology.*

Proof Obviously, 3 implies 2; the converse is also easy to check directly, but we won't need it because we'll prove $1 \longrightarrow 3$ and $2 \longrightarrow 1$.
Proof of $1 \longrightarrow 3$: Assume ϕ is usHv, and consider the first-order instance ψ obtained by replacing each propositional variable p by a formula of the form $\forall x \, \exists y \, P(x,y)$, using different binary predicate symbols

P for different propositional variables p. By assumption, ψ is simply Herbrand valid, so let θ be a tautologous closed instance of ψ_H. Thus, θ is like ϕ except that each occurrence of a propositional variable p has become an occurrence of $P(t,u)$ for some closed terms t and u — possibly different closed terms for different occurrences of the same p.

In fact, we rather rarely get the same $P(t,u)$ repeated. To see this (and to make "rather rarely" precise), recall the first step in the definition of the Herbrand form ψ_H: Rename bound variables so that no variable is quantified twice. Thus, if we consider two positive occurrences of p in ϕ, the corresponding subformulas of ψ will, after this renaming, look like $\forall x\,\exists y\,P(x,y)$ and $\forall x'\,\exists y'\,P(x',y')$, and in the Herbrand form ψ_H these will have become $P(X,y)$ and $P(X',y')$, with different constant symbols X and X' in the first argument place. So these two positive occurrences of p in ϕ will become different atomic formulas in θ.

Similarly, if we consider two negative occurrences of a propositional variable p in ϕ, they become $P(x,Y(x))$ and $P(x',Y'(x'))$ in ψ_H, with different function symbols Y and Y'. So the corresponding atomic formulas in θ are different.

Thus, the only way two atomic formulas in θ can be the same is for one to arise from a positive occurrence and the other from a negative occurrence of some p in ϕ. In particular, no atomic formula can occur three times in θ. If we regard the atomic formulas occurring in θ as propositional variables, then the preceding discussion shows that θ is a binary tautology. Since ϕ is clearly a variable-merging instance of it, obtained by substituting the original variables p for all atomic formulas in θ that begin with the corresponding predicate symbols P, we have established 3.

Proof of $2 \longrightarrow 1$: We must show that any first-order instance of any propositional instance of a binary tautology is simply Herbrand valid. The proof proceeds by three reductions of the problem, after which the remaining work is quite easy. The first reduction is to observe that, since "instances of instances are instances" we need only show that any first-order instance ψ of any binary tautology θ is simply Herbrand valid.

The second reduction is to arrange that, without loss of generality, whenever a propositional variable occurs twice in θ, one occurrence is positive and the other negative. Indeed, suppose p had two positive occurrences, and let θ' be obtained from θ by replacing one of these occurrences by a new propositional variable p'. Clearly, ψ is a first-order instance of θ' and θ' is binary. Furthermore, θ' is a tautology. To see this, suppose we had a truth assignment falsifying θ'. It must give p and p' different truth values, for otherwise it would also falsify θ. But then if we change the value "true" of p or of p' to "false," the new truth assignment will still falsify θ' because both p and p' occurred positively. So the new truth assignment would falsify the tautology θ. This contradiction shows that we can eliminate a repetition of a variable when both its occurrences are positive. An analogous argument eliminates repetitions when both occurrences are negative.

Before proceeding to the third reduction, we summarize the present situation. We assume from now on that θ is a tautology in which each propositional variable has at most one positive occurrence and at most one negative occurrence. Let ψ be any first-order instance of θ, and let ψ_H be its Herbrand form. We shall complete the proof by finding a closed instance of ψ_H that is, when regarded as a propositional formula (viewing its atomic subformulas as propositional variables) a propositional instance of θ and therefore a tautology.

Each propositional variable p in θ becomes some (first-order) subformula π of ψ, which in turn becomes (at worst) two subformulas π^+ and π^- in ψ_H. Here π^+ occurs in place of the positive occurrence of p and π^- occurs in place of the negative occurrence of p in θ. (If p had only one occurrence in θ, then one of $\pi^\pm$ is absent.) Notice that, unless π happened to be quantifier-free, π^+ and π^- are different formulas, for any quantifier in π will be universal in one of the two copies of π and existential in the other. (Remember that negations always get pushed in past the quantifiers.) Our task is to form a closed instance of ψ_H in such a way that the corresponding instances of π^+ and π^- are identical. Then this instance of ψ_H will be an instance of θ (obtained by replacing each propositional variable p in θ by the common instance of the corresponding formulas $\pi^\pm$), so it will be a tautology, and the proof will be complete.

The third reduction is to see that, when defining this instance of ψ_H, we may focus on a single p and the two formulas $\pi^\pm$ that it developed into in going from θ to ψ_H. (If p had only one occurrence in θ, then we need not concern ourselves with it; our only task is to make sure that, when π^+ and π^- both exist, then our closed instantiation makes them identical.) While focusing on one p, we shall define the closed terms that are to be substituted for the free variables in π^+ and π^-; we define them in such a way as to make the resulting instances of π^+ and π^- agree. What we do here for a particular p will not interfere with the corresponding efforts on behalf of other propositional variables q in θ, because those efforts will define what is to be substituted for *different* variables — different because of the renaming step in the definition of Herbrand form. This completes the three reductions.

Concentrating therefore on a particular p, assume that it occurs twice in θ, once positively and once negatively, and assume that its replacement π in ψ has no variable bound twice. (It would be rewritten this way in forming ψ_H, so there is no harm in supposing it is already written this way.) Let $x_1, x_2, \ldots, x_n$ be the variables occurring in the sentence π, ordered in such a way that, if the quantifier of x_i lies in the scope of the quantifier of x_j then $j < i$. For example, the variables could be ordered according to the left-to-right order of occurrence of their quantifiers in π. Notice that, since p occurs twice in θ and π correspondingly occurs twice in ψ, each x_i will become two distinct variables in the construction of ψ_H (see step 1 in the definition of Herbrand form). One of these will be universally quantified, the other existentially. The former becomes, in ψ_H, a function symbol which we call X_i; the latter remains a variable, which we call x'_i. The argument places of X_i are occupied by some x'_j's with smaller subscripts. We write $X_i(x'_1, x'_2, \ldots, x'_{i-1})$, but with the understanding that not all of the indicated arguments have to be present.

Now we define, by induction on i, the closed term t_i to be substituted for the variable x'_i in forming the desired instance of ψ_H. For any particular i, assume by induction that t_j is already defined for all $j < i$. Define t_i to be $X_i(t_1, t_2, \ldots, t_{i-1})$, with the notation as described at the end of the last paragraph (so some of the earlier t_j's may not really be present). It is easy to verify that this definition makes the resulting instances of π^+ and π^- identical. So the proof is complete. □

Remark 5.3 We are now in a position to be more explcit about the difference, already mentioned in Remark 3.2, between the two common definitions of the biconditional,

- $(A \longrightarrow B) \wedge (B \longrightarrow A)$, i.e., $(\neg A \vee B) \wedge (\neg B \vee A)$, and
- $(A \wedge B) \vee (\neg A \wedge \neg B)$.

With the first definition, $A \longleftrightarrow A$ is usHv, because it is an instance of the binary tautology $(\neg A \vee A) \wedge (\neg B \vee B)$. With the second definition, $A \longleftrightarrow A$ is $(A \wedge A) \vee (\neg A \wedge \neg A)$, which is easily seen not to be an instance of a binary tautology. Thus, the difference between the two definitions of the biconditional is essential in the context of (universal) simple Herbrand validity.

6 Modus Ponens

In the introduction, we mentioned that simple Herbrand validity is not well-behaved and therefore serves mainly as a step on the way to the concept of universal simple Herbrand validity. In this section, we indicate the sort of misbehavior that we had in mind. We give an example of two simply Herbrand valid sentences, of the forms ϕ and $\phi \longrightarrow \psi$, such that ψ is not simply Herbrand valid. That is, the rule of modus ponens is not sound for simple Herbrand validity. We shall see later that it is sound for universal simple Herbrand validity.

We work in a vocabulary with one unary predicate symbol M and, in accordance with Convention 3.1, one constant symbol c. Let ϕ be the sentence

$$\exists u\, M(u) \longrightarrow \exists v\, M(v),$$

and let ψ be

$$\exists x\, \forall y\, (M(y) \longrightarrow M(x)).$$

Then the Herbrand form of ϕ is $M(U) \longrightarrow M(v)$, which has a tautologous closed instance $M(U) \longrightarrow M(U)$.

The Herbrand form of $\phi \longrightarrow \psi$ is

$$(M(u) \longrightarrow M(V)) \longrightarrow (M(Y(x)) \longrightarrow M(x)).$$

This also has a tautologous closed instance, obtained by replacing x with V and u with $Y(V)$.

The Herbrand form of ψ is $M(Y(x)) \longrightarrow M(x)$. In any closed instance of this, the antecedent is different from the consequent, so no single closed instance of ψ_H is a tautology.

Thus, both ϕ and $\phi \longrightarrow \psi$ are simply Herbrand valid but ψ is not.

In this example, although no single closed instance of ψ_H is a tautology, there are two instances whose disjunction is a tautology. Replace x by c for one instance and by $Y(c)$ for the other.

The example can be modified so that ψ is much farther from being simply Herbrand valid. Given any positive integer s, work in a vocabulary having s unary predicate symbols $M_1, \ldots, M_s$ and one constant symbol, let ϕ be

$$(\exists u_1\, M_1(u_1) \longrightarrow \exists v_1\, M_1(v_1)) \wedge \cdots \wedge (\exists u_s\, M_s(u_s) \longrightarrow \exists v_s\, M_s(v_s)),$$

and let ψ be

$$\exists x_1 \ldots \exists x_s\, \forall y\, [\neg M_1(y) \vee (M_1(x_1) \wedge \neg M_2(y)) \vee \ldots \\ \cdots \vee (M_{s-1}(x_{s-1}) \wedge \neg M_s(y)) \vee M_s(x_s)].$$

Then both ϕ and $\phi \longrightarrow \psi$ are simply Herbrand valid, but a tautologous disjunction of closed instances of ψ_H must use at least $s+1$ instances.

7 Connection with Affine Logic

In (2), I proposed a game semantics for Girard's affine logic (3) and proved a soundness theorem for it. For the additive fragment, there is a completeness theorem also (2, Section 4), but this is not so for the multiplicative fragment. The game-valid multiplicative formulas are a strictly wider class than the ones provable in affine logic. The former class admits, however, a syntactic characterization, given by the "Multiplicative Validity Theorem" of (2, Section 5): Up to notational differences (the $A^{\perp}$ and $A \otimes B$ of affine logic corresponding to the $\neg A$ and $A \wedge B$ of traditional logic), the game-valid multiplicative formulas are exactly the propositional instances of binary tautologies. That is, they are exactly the usHv formulas.

There is some empirical evidence for the naturalness of the game semantics of (2). First, the operations on games used to interpret the connectives arose originally in purely game-theoretic considerations (1). Second, they arose independently in Japaridze's analysis (5) of "effective truth." The games and strategies considered in (5) are quite different from those in (2), for in Japaridze's games all plays are finite and strategies are required to be effectively computable. Nevertheless, Japaridze finds the same notion of validity for the multiplicative fragment.

That this same notion of validity has now arisen a third time, from the analysis of Herbrand disjunctions, increases the evidence for the importance of these concepts. It suggests that instances of binary tautologies occur naturally in attempts to analyze resource consciousness and deserve more attention in their own right. I am aware of only one paper, Jaśkowski's (6), in which binary tautologies are studied for their own sake. They occur there as the solution to the problem of characterizing the provable formulas of a certain deductive system.

It is well known that the new function symbols occurring in the Herbrand form ϕ_H of a sentence ϕ can be interpreted as (part of) a strategy for player $\forall$ in a game where $\exists$ tries to confirm the truth of ϕ while $\forall$ tries to refute ϕ. This might lead one to regard a connection between Herbrand's theorem and game semantics as unsurprising. I know of no way, however, to make these ideas precise. In particular, I do not know how to use the relatively easy proof of Theorem 5.2 in the present paper to simplify the considerably harder proof of the multiplicative validity theorem in (2).

Finally, it should be pointed out that, in contrast to what we showed in Section 6 for simple Herbrand validity, universal simple Herbrand validity is preserved by modus ponens and by the slightly more general cut rule. A proof using too much machinery is to apply the soundness theorem of (2) which says that cut (along with all the other rules of affine logic) is sound for game-validity, and to invoke the fact that game-validity in the multiplicative fragment coincides with universal simple Herbrand validity. An alternate proof is to show by a direct combinatorial argument that cut preserves the property of being an instance of a binary tautology; that combinatorial argument is left as an amusing exercise for the reader (as it was in (2)).

References

[1] A. Blass, "Degrees of indeterminacy of games," *Fund. Math.* 77 (1972) 151–166.

[2] A. Blass, "A game semantics for linear logic," *Ann. Pure Appl. Logic* 56 (1992) 183–220.

[3] J.-Y. Girard, "Linear logic," *Theoret. Comput. Sci.* 50 (1987) 1–102.

[4] J. Herbrand, *Recherches sur la théorie de la démonstration*, Travaux de la Société des Sciences et Lettres de Varsovie, Classe III, sciences mathématiques et physiques, no. 33 (1930).

[5] G. Japaridze, "A constructive game semantics for the language of linear logic," *Ann. Pure Appl. Logic* 85 (1997) 87–156.

[6] S. Jaśkowski, "Über Tautologien, in welchen keine Variable mehr als zweimal vorkommt," *Z. Math. Logik Grundlag. Math.* 9 (1963) 219–228.

[7] J. R. Shoenfield, *Mathematical Logic*, Addison-Wesley (1967).