Discussion Paper No. 1066R

ORDER INDEPENDENCE FOR ITERATED WEAK DOMINANCE

by

Leslie M. Marx
Jeroen M. Swinkels

Northwestern University

September 10, 1996

Order Independence for Iterated Weak Dominance*

LESLIE M. MARX
W.E. Simon Graduate School of Business Administration,
University of Rochester, Rochester, New York 14627

AND

JEROEN M. SWINKELS
J.L. Kellogg Graduate School of Management,
Northwestern University, Evanston, Illinois 60208

September 10, 1996

Abstract

In general, the result of the elimination of weakly dominated strategies depends on order. We define *nice weak dominance*. Under nice weak dominance, order does not matter. We identify an important class of games under which nice weak dominance and weak dominance are equivalent, and so order under weak dominance does not matter. For all games, the result of iterative nice weak dominance is an upper bound on the result from any order of weak dominance. The results strengthen the intuitive relationship between backward induction and weak dominance, and shed light on some computational problems relating to weak dominance. *Journal of Economic Literature* Classification Number: C72.

*Marx gratefully acknowledges the support of an NSF Graduate Fellowship and a Sloan Doctoral Dissertation Fellowship. Swinkels gratefully acknowledges support from the National Science Foundation. The authors are grateful to two referees, an associate editor, Paul Milgrom, Roger Myerson, and Phil Reny for useful comments and discussions.

I. Introduction

As is well known, the result of the iterative removal of weakly dominated strategies can depend on the order of removal.[1] Consider G1 with $x \geq 1$.

		II	
		L	R
I	T	2,3	x,3
	M	1,0	0,1
	B	0,1	1,0

Fig 1. G1.

Depending on the order of elimination, the set of strategies that remains after iterative removal of weakly dominated strategies can be $\{T, L\}$, $\{T, R\}$, or $\{T, L, R\}$.

In this paper, we define the *transference of decision maker indifference* (TDI) condition. TDI is satisfied by the normal form of any extensive form for generic assignment of payoffs to terminal nodes and by some important games that do not admit generic extensive forms, including discretized versions of first price auctions. We show that under TDI any two games obtained from the original by the iterative elimination of weakly dominated strategies (subject to no more eliminations being possible) are strategically equivalent. That is, the two games differ only by the addition or removal of redundant strategies and a renaming of strategies. In G1, TDI is satisfied if and only if $x = 2$, and if $x = 2$, then $\{T, L\}$, $\{T, R\}$, and $\{T, L, R\}$ are all strategically equivalent "reductions" of G1.

We derive our results in terms of a concept we term *nice weak dominance*. Essentially, a weak dominance is nice if TDI is satisfied for the strategies involved in that weak dominance, and so in particular, if a game satisfies TDI, then weak dominance and nice weak dominance are identical.

We show that regardless of the game, the result of iterative removal by nice weak dominance does not depend on order. So for games satisfying TDI, the result of weak dominance is also order independent. And, in games not satisfying TDI, we show that the result of weak dominance is always essentially a subset of the result of nice weak dominance. So, even for games where TDI fails, the result allows us to easily identify strategies that must be removed regardless of order, and to establish an upper bound on the amount by which the results of different orders of elimination can differ.

We begin by discussing these results for the case of weak dominance involving pure strategies. However, the results translate fairly directly into the world of mixed strategies. We discuss this in Section V.

---

[1] Let s_i and r_i be strategies for player i. Given opponents' strategy set W_{-i}, s_i weakly dominates r_i if $\pi_i(s_i, t_{-i}) \geq \pi_i(r_i, t_{-i})\ \forall t_{-i} \in W_{-i}$ and $\pi_i(s_i, t_{-i}) > \pi_i(r_i, t_{-i})$ for some $t_{-i} \in W_{-i}$, where π_i is player i's payoff function.

Backward induction and weak dominance seem intimately related, both in motivation and in that every sequence of removals of actions by backward induction is echoed by a sequence of removals of strategies by weak dominance. An uncomfortable feature of this relationship is that backward induction has a deterministic outcome, while weak dominance does not. We show that eliminations of actions under backward induction are in fact eliminations by nice weak dominance. Thus the order independence of backward induction and the order dependence of iterated weak dominance are related to the fairly simple difference between weak dominance and nice weak dominance, strengthening the intuitive connection between backward induction in the extensive form and weak dominance in the normal form.

Finally, we use our result to make a comment on the work of Gilboa *et al.* (1993) concerning the complexity of iterative weak dominance.

II. THE LITERATURE

A number of previous papers have explored the issue of order independence. Gilboa, Kalai, and Zemel (1990) give conditions on a dominance operator that are sufficient for the order of elimination not to matter. These conditions are satisfied by strong dominance, but not by weak dominance (or nice weak dominance).[2]

Rochet (1980) considers the following condition:

$$\pi_i(s) = \pi_i(t) \Rightarrow \pi_j(s) = \pi_j(t) \text{ for all } i, j \in N, \ s, t \in \Pi_{i \in N} S_i, \tag{1}$$

where N is the set of players, S_i is the set of strategies for player i, and π_i is player i's payoff function (see Section IV for formal definitions).

Rochet shows that if a game satisfying (1) is dominance solvable (when one eliminates all weakly dominated strategies at every stage), then the same outcome is obtained regardless of the order of elimination of weakly dominated strategies. Furthermore, Rochet shows that any normal form game derived from an extensive form

[2]Gilboa *et al.* also consider a version of weak dominance which does not require the strict inequality. We shall refer to this as *very weak dominance*. They claim, but do not prove, that with some additional conditions (which are satisfied by both weak and very weak dominance), arguments similar to theirs can be used to show that order of removal under a relationship dom will not matter if

$$[x \text{ dom } y \text{ and } y \text{ dom } x] \Rightarrow [x \text{ and } y \text{ are payoff equivalent for all players}],$$

a condition which is similar to our TDI. This claim is false as stated, because for weak dominance the antecedent never holds (since it can never be the case that x and y each weakly dominate the other) and so the condition is vacuously satisfied, but order can clearly matter. The claimed result is correct as stated for very weak dominance, and many of the key ideas in their analysis reappear in the course of our proof. Gilboa *et al.* note that their condition is satisfied for zero sum games under very weak dominance.

of perfect information satisfying the extensive form analogy to (1) (i.e., satisfying that one player is indifferent over two terminal nodes only if all players are) is dominance solvable with the same outcome as that determined by backward induction on the extensive form. Rochet shows by example that this need not be the case when the extensive form does not satisfy this condition.[3]

Gretlein (1983) works with games in which each player's preferences over the set of possible outcomes (i.e. payoff vectors) is strict. In such games, Gretlein shows that the set of outcomes that results from iterated elimination of weakly dominated strategies (subject to no more eliminations being possible) is the same regardless of the order of eliminations. This condition implies (1).[4]

Our condition, TDI, is weaker than (1). In Section III we give examples of interesting games satisfying TDI but not (1). Both Rochet and Gretlein consider only domination by pure strategies, while we extend our results to mixed strategies. In addition, neither Rochet nor Gretlein has anything to say about games for which (1) is not satisfied, but our notion of nice weak dominance allows us to establish a bound on the outcome of iterated weak dominance even for games not satisfying TDI.

III. Transference of Decision Maker Indifference

The normal form of a generic extensive form game always satisfies condition (1): for $\pi_i(s) = \pi_i(t)$ to hold, it must be that s and t reach the same terminal node. But then $\pi_j(s) = \pi_j(t)$ for all $j \in N$. Moulin (1979) identifies a class of voting games satisfying (1). So, even under condition (1), establishing that the result of iterative elimination under weak dominance does not depend on order is of considerable use. However, there are important games not satisfying (1) to which we would also like our results to apply.

As an example, consider a first price auction. To make this a finite strategy game, assume that players receive signals about the value of the object that are drawn from a finite set Ω (the analysis to follow does not depend on the manner in which

[3]See also Moulin (1986, Chapter 4.2) on Rochet's robustness result, dominance solvability (using weak dominance), and the relationship between (1) and extensive-form games. Moulin (1984) gives conditions for a game to be dominance solvable (using weak dominance) and shows that for a certain class of games, dominance solvability implies Cournot stability.

[4]Also, note that Gretlein's result states only that the set of outcomes is the same, not that the games that remain after iterated removal are in any sense equivalent. If one is interested, for example, in Nash equilibria of the game which results from the iterative removal, then it is important to know that different orders cannot yield games which differ as the following two games do:

	2	
1	6,2	2,6
	1,1	4,0

	2	
1	6,2	1,1
	4,0	2,6

signals are related across players), and that players are restricted to make bids that are integer multiples of a penny up to some (large) maximum. Then, each player's strategy space is the set of all maps from signals in Ω to allowable bids, and so is finite. Let r_i and s_i be two strategies for some player i that differ only in that for some signal $\omega \in \Omega$, $r_i(\omega) < s_i(\omega)$; i.e., let r_i and s_i differ only in that for some signal, r_i specifies a smaller bid than s_i. Consider any strategy profile t_{-i} for the other players such that the largest bid possible under t_{-i} is less than $r_i(\omega)$. Then, $\pi_{-i}(r_i, t_{-i}) = \pi_{-i}(s_i, t_{-i})$. To see this, note that when i's signal is not ω, his behavior is the same under r_i and s_i. When i's signal is ω, players other than i do not receive the object, and so they receive a payoff of 0 in either case. However, $\pi_i(r_i, t_{-i}) \neq \pi_i(s_i, t_{-i})$. Thus, this game does not satisfy (1).

The condition failed in this example because, once a player has lost, he is indifferent over the amount by which he loses. For almost all specifications of how signals map to valuations, this is the only way in which i can be indifferent between pure strategy profiles (r_i, t_{-i}) and (s_i, t_{-i}) — by having those two strategy profiles differ only in how much i loses by when i loses. We formalize this in Appendix A. So, while a unilateral change of pure strategy by player i can change his payoff while leaving his opponents indifferent, the opposite cannot occur: i cannot be indifferent about this unilateral change while affecting his opponents' payoffs. Thus, while the game does not satisfy (1), it does satisfy the following weaker condition

$$\begin{gathered} \pi_i(r_i, s_{-i}) = \pi_i(t_i, s_{-i}) \Rightarrow \pi_j(r_i, s_{-i}) = \pi_j(t_i, s_{-i}) \\ \text{for all } i, j \in N,\ r_i, t_i \in S_i,\ s_{-i} \in S_{-i}. \end{gathered} \tag{2}$$

This is weaker than (1), because the agreement of player i's payoffs across two different strategy profiles only implies agreement for the other players if the strategy profiles differ only by the action of player i. So in game G2, (1) implies $b = c = d$, while (2) implies $b = c$. Note in particular that while (B,L) and (B,R) differ only in the action of one player, it is the payoffs of the non-decision maker that agree, and so (2) has no power.

		II: L	II: R
I	T	a,b	x
	B	a,c	a,d

Fig 2. G2.

We shall refer to (2) as the *transference of decision maker indifference* (TDI) condition: whenever the decision maker is indifferent between two profiles that differ only in her action, that indifference is transferred to the other players as well.[5] We

[5]Interestingly, TDI is formally equivalent to the "nonbossiness" condition in social choice theory. See for example Satterthwaite and Sonnenschein (1981).

will show for any finite player game satisfying TDI that any two games that are achieved by the iterative removal of weakly dominated strategies (subject to no more removals being possible) are equivalent.

For what types of games is (2) satisfied, but (1) not? We have already seen the example of the first price auction. However, the class of games that satisfy (2) but not (1) is much broader. Consider any game in which at some point some player has the option to quit the game and collect an outside option. The game then continues with a subgame played between the remaining players. Examples include patent races and oligopoly with an endogenous number of firms. In such games, the payoff of the exiting player is independent of the outcome of the subgame, but the payoffs of other players are not. So, (1) is violated. However, because the exiting player makes no decisions in the subgame, these payoff ties do not violate (2). And, if the player making the exit decision is not indifferent over the various times at which he might exit, then the game as a whole will satisfy (2). In the case of the patent race for example, the exiting player presumably incurs costs as long as he is in the race, and so he is not indifferent between different exit times.

As another example, suppose a particular public good will be provided if at least k players vote that it should be purchased. If purchased, the cost of the public good will be shared equally among the people who vote for it. For generic valuations of the public good, if player i is indifferent between voting for the public good and not, it must be that the good is not purchased in either case, so everyone is indifferent. However, (1) is clearly not satisfied.

Consider students competing for grades. The top 30% of the class receive A's, the next 60% B's, and the remainder C's. Students simultaneously choose actions (effort levels, attendance, study strategies) from a finite set. A student's payoff is a function of her actions and grade. Then, if one student changes her action and moves from one grade level to another, the grades of (at most) two other students will be affected, and so (1) will fail. However, for generic assignment of payoffs to (action, grade) combinations, a student is never indifferent over two strategy profiles that differ only in her actions, so (2) will be satisfied. Other examples of this type of "placement game" might include competition by workers for promotions or advertising firms for accounts. Notice that these games bear some similarities to first price auctions, and that indeed, first price auctions are an example of this class.

Next, consider "group formation" games. In these games, there is an initial stage during which the players are assigned to groups, modeled by a process in which each player chooses an action and then, as a function of the vector of actions, is assigned to a group. Once the grouping process is finished, the members of each group play a game among themselves. Payoffs depend only on the members of one's group and the strategy profile chosen within that group. If the group formation process is of interest to the modeler, then all these "subgames" must be incorporated into a single extensive form. This will involve violating (1) because a member of one group will be indifferent over the actions of players in other groups. Thus, there will be many pairs of strategy profiles where some but not all players have equal payoffs. However, the game can easily satisfy (2).[6]

Another interesting example is signaling games in which the payoff of the receiver is independent of the signal of the sender. Then, for any fixed action by the receiver and type of the sender, the payoff of the receiver is unchanging as the signal changes, but if signals have different costs, the payoff of the sender is not. Thus (1) is violated, but (2) holds generically. An example in this spirit is "burn the dollar" games, in which the payoff of the non-burner is dependent only on the outcome in the underlying game, while the payoff of the burner depends both on the outcome in the underlying game and on whether or not the dollar was burned, so (1) could not be satisfied, but for generic payoffs in the underlying game, (2) is.

There are interesting games for which (2) fails. In a second price auction, for example, a player is indifferent between non-winning bids, but the winner may well not be.

IV. Formalities and the Main Result

We work with finite strategy, finite player, normal form games. Players $i \in N \equiv \{1, ..., n\}$ have finite pure strategy spaces S_i. Payoffs are given by $\pi : \Pi_{i \in N} S_i \rightarrow \Re^n$. The payoff function π is extended to mixed strategies in the standard way. We assume, without loss of generality, that $S_i \cap S_j = \emptyset$ for all $i, j \in N$, $i \neq j$. So, without ambiguity, we can drop the player subscripts on the strategy names. Let $S \equiv \bigcup_{i \in N} S_i$. For $W \subseteq S$, let the strategies in W that belong to i be denoted by $W_i \equiv W \cap S_i$. Say that $W \subseteq S$ is a *restriction* of S if $\forall i, W_i \neq \emptyset$. Note that any restriction W of S generates a unique game given by strategy spaces W_i and the restriction of π to $\Pi_{i \in N} W_i$. We will denote this game by (W, π). We similarly define $W_{-i} \equiv \Pi_{j \neq i} W_j$. A typical element $x_{-i} \in W_{-i}$ thus specifies a strategy $x_j \in W_j$ for each $j \neq i$.

Definition 1. Let W be a restriction of S, and let $r_i, s_i \in S_i$. Then

(i) r_i *very weakly dominates* s_i on W, written $r_i \; VWD_W \; s_i$, if $\pi_i(r_i, x_{-i}) \geq \pi_i(s_i, x_{-i}) \; \forall x_{-i} \in W_{-i}$, and

[6] Assume for example that (a) no player is indifferent between any two outcomes which occur in two different groupings, (b) if a player changes his action in the grouping phase and this affects the group to which anyone belongs, then it also affects the group to which the player himself belongs, (c) players observe the group to which they are assigned and (of course) recall their own action in the matching stage, but nothing else, and (d) each subgame satisfies TDI. Then, consider any pair of strategy profiles which differ only in the actions of some player i. If player i is indifferent between the two outcomes, then by (a) and (b) he cannot have affected the groupings. So, by (c) any player not in i's group will get the same payoff. Finally, because of (d), any change in the outcome in i's subgame will be consistent with TDI.

(ii) r_i *weakly dominates* s_i on W, written $r_i \; WD_W \; s_i$, if $r_i \; VWD_W \; s_i$ and, in addition, $\pi_i(r_i, z_{-i}) > \pi_i(s_i, z_{-i})$ for some $z_{-i} \in W_{-i}$.

This definition is phrased entirely in terms of pure strategies. Of course, when determining if s_i is weakly dominated on W, it makes no difference whether one considers W_{-i} or mixtures over W_{-i}. However, it is possible that s_i is weakly dominated by a mixture over W_i but not by any element of W_i. We extend our results to weak dominance by mixtures in Section V.

DEFINITION 2. A weak dominance or very weak dominance involving r_i and s_i on W is *nice* if for all $x_{-i} \in W_{-i}$, if $\pi_i(r_i, x_{-i}) = \pi_i(s_i, x_{-i})$ then $\pi(r_i, x_{-i}) = \pi(s_i, x_{-i})$.

That is, a dominance is nice if TDI is satisfied with respect to the strategies involved. Write $r_i \; NWD_W \; s_i$ if r_i weakly dominates s_i on W and the weak dominance is nice. Say that s_i is *nicely weakly dominated on* W if there is r_i in $W \backslash s_i$ with $r_i \; NWD_W \; s_i$. Define $NVWD_W$ and *nicely very weakly dominated on* W similarly for very weak dominance.

DEFINITION 3. Let V be a restriction of S, and let W be a restriction of V. Then W is a *reduction of* V *by (nice) (very) weak dominance* if $W = V \backslash X^1, \ldots, X^m$, where $\forall k, X^k \subset S$, and $\forall x \in X^k$, $\exists z \in V \backslash X^1, \ldots, X^k$ such that z (nicely) (very) weakly dominates x on $V \backslash X^1, \ldots, X^{k-1}$. W is a *full reduction of* V *by (nice) (very) weak dominance* if W is a reduction of V by (nice) (very) weak dominance and no strategies in W are (nicely) (very) weakly dominated on W.

That is, a reduction is the result of iteratively removing sets of strategies that are dominated in the appropriate sense, and a full reduction is one in which no dominances of the appropriate sort remain.

If each of the sets X^i in the previous definition are singletons, then W is a *one-at-a-time reduction* of V (by whatever dominance relationship is under discussion).

Clearly NWD is more restrictive than WD and $NVWD$, which in turn are each more restrictive than VWD. So, any (one-at-a-time) reduction by NWD is also a (one-at-a-time) reduction by WD and $NVWD$, and any (one-at-a-time) reduction by WD or $NVWD$ is a (one-at-a-time) reduction by VWD.

LEMMA A. *Let* W *be a reduction of* V *by (nice) very weak dominance. Then* W *is a one-at-a-time reduction of* V *by (nice) very weak dominance.*

Proof. Clearly (nice) very weak dominance on W implies (nice) very weak dominance on any subset of W. So, the set of strategies removed at each step can instead be removed one at a time in any order. A simple induction then yields the result. ■

DEFINITION 4. Let V and W be restrictions of S. V is *equivalent to a subset* of W if there exist one-to-one maps $m_i : V_i \to W_i$, $i \in N$, such that $\pi(x) = \pi(m_1(x_1), \ldots, m_n(x_n))\ \forall x \in V$.

So, V differs from a subset of W (specifically $m(V)$) only by a renaming.

OBSERVATION 1. The relation "equivalent to a subset" is transitive.

DEFINITION 5. Let W be a restriction of S, and let $r_i, s_i \in S_i$. Then r_i is *redundant* to s_i on W, written $r_i\ R_W\ s_i$, if $\pi(r_i, x_{-i}) = \pi(s_i, x_{-i})\ \forall x_{-i} \in W_{-i}$. A strategy s_i is *redundant* on W if there is $r_i \in W_i \backslash s_i$ with $r_i\ R_W\ s_i$.

OBSERVATION 2. $r_i\ NVWD_W\ s_i$ if and only if either $r_i\ NWD_W\ s_i$ or $r_i\ R_W\ s_i$.

Clearly VWD is transitive. In addition:

OBSERVATION 3. $NVWD$ is transitive: i.e., $r_i\ NVWD_W\ s_i$ and $s_i\ NVWD_W\ t_i$ imply $r_i\ NVWD_W\ t_i$.

To see this, suppose $r_i\ NVWD_W\ s_i$ and $s_i NVWD_W\ t_i$, and consider any $s_{-i} \in W_{-i}$ such that $\pi_i(r_i, s_{-i}) = \pi_i(t_i, s_{-i})$. Then $\pi_i(r_i, s_{-i}) = \pi_i(s_i, s_{-i}) = \pi_i(t_i, s_{-i})$, and so $\pi(r_i, s_{-i}) = \pi(s_i, s_{-i}) = \pi(t_i, s_{-i})$. Thus, $r_i\ NVWD_W\ t_i$.

DEFINITION 6. Let W be a restriction of S. A set V is said to be a *reduction of W by (nice) weak dominance/redundance/substitution* if V can be obtained from W by letting $V^0 = W$ and performing a finite number of iterations of the following process: obtain V^{j+1} from V^j by either (i) letting $V^{j+1} = V^j \backslash s^j$, where s^j is (nicely) weakly dominated on V^j by an element of $V^j \backslash s^j$, or (ii) letting $V^{j+1} = V^j \backslash s^j$, where s^j is redundant on V^j to an element of $V^j \backslash s^j$, or (iii) letting $V^{j+1} = (V^j \backslash s^j) \cup r$, where $r \in S$ and r is redundant on V^j to s^j.

Note that if W is a restriction of S, then any reduction of W by (nice) weak dominance/redundance/substitution is also a restriction of S.

LEMMA B. *Let W be a restriction of S, let T be a reduction of W by (nice) very weak dominance/redundance/substitution, and let $s \in W$. Then there exists $s' \in T$ such that s' (nicely) very weakly dominates s on T.*

Proof. Let $s \in W$. Since T is obtained from W by (nice) very weak dominance/redundance/substitution, there exist restrictions $T^0, T^1, \ldots, T^m$ such that $T^0 = W$ and $T^m = T$ and each T^{j+1} is obtained from T^j as in the definition of a reduction by dominance/redundance/substitution. The Lemma is satisfied for T^0 because $s \in T^0$.

so s itself serves. Assume that the Lemma holds for T^j. Then there exists $s' \in T^j$ such that s' (nicely) very weakly dominates s on T^j. If $s' \in T^{j+1}$, then s' serves and the Lemma holds for T^{j+1}. Note in particular that if a strategy for another player is removed in going from T^j to T^{j+1}, then since the requisite inequalities and equalities hold on T^j, they continue to hold on any subset, and in particular on T^{j+1}, while if a strategy is substituted, then the requisite relations continue to hold by the definition of redundance.

Suppose that $s' \notin T^{j+1}$. Then s' is the strategy that was eliminated when T^{j+1} was obtained from T^j. If s' was eliminated by (nice) weak dominance or redundance by some $t \in T^j \backslash s' = T^{j+1}$, then let $s'' = t$. If s' was eliminated and replaced by redundant strategy $r \in W \backslash s'$, then let $s'' = r$. In either case, $s'' \in T^{j+1}$ and s'' (nicely) very weakly dominates s on T^{j+1}. The result follows by induction. ■

Lemma C is the key to the results. Roughly, it says that if you can get from W to $\hat{W}$ by $NVWD$, and you can get from W to V by (nice) very weak dominance, then you can get from V to what is essentially a subset of $\hat{W}$ by iteratively removing strategies that are either redundant or (nicely) weakly dominated.

LEMMA C. *Let W be a restriction of S, let $\hat{W}$ be a reduction of W by nice very weak dominance, and let V be a reduction of W by (nice) very weak dominance. Then there exists $\hat{V}$ equivalent to a subset of $\hat{W}$, where $\hat{V}$ is obtainable from V by the iterative removal of strategies that are either (nicely) weakly dominated or redundant.*

Proof. Since $\hat{W}$ is a reduction of W by nice very weak dominance, by Lemma A we can write $\hat{W} = W \backslash x^1, ..., x^m$, where for $k = 1, ..., m$, $W^k = W \backslash x^1, ..., x^k$ is a (one-at-a-time) reduction of $W \backslash x^1, ..., x^{k-1}$ by nice very weak dominance. Let $W^0 = W$.

We proceed by induction. For some $j \in \{0, ..., k-1\}$ assume that we have V^j and m^j such that

(1) V^j can be obtained from V by iteratively eliminating strategies that are either (nicely) weakly dominated or redundant,

(2) m^j is a one to one map such that $\forall i \in N$, $\forall s_i \in V_i^j$, $m^j(s_i) \in W_i^j$ and $\forall t \in \prod_{i \in N} V_i^j$, $\pi(t) = \pi(m^j(t_1), ..., m^j(t_n))$, and

(3) $m^j(V^j)$ is a reduction of V by (nice) weak dominance/redundance/substitution.

We will exhibit V^{j+1} and m^{j+1} satisfying (1) – (3). Since (1) – (3) are trivially satisfied for $j = 0$ by taking $V^0 = V$ and m^0 as the identity map on V^0, this will establish the Lemma.

So, given V^j and m^j, let us construct V^{j+1} and m^{j+1}.

If $m^j(V^j) \subseteq W^{j+1}$, let $V^{j+1} = V^j$ and $m^{j+1} = m^j$.

Suppose $m^j(V^j) \not\subseteq W^{j+1}$. Then since $m^j(V^j) \subseteq W^j$, $x^{j+1} \in m^j(V^j)$. Let i be the player to whom x^{j+1} belongs and let $r_i \in W_i^{j+1}$ be a strategy that nicely very weakly dominates x^{j+1} on W^j. Since $m^j(V^j)$ is a reduction of V by (nice) weak dominance/redundance/substitution, and since V is a reduction of W by (nice) very weak dominance. Lemma B implies that there is $y \in m^j(V^j)$ that (nicely) very weakly dominates r_i, and thus x^{j+1}, on $m^j(V^j)$. Then, either y (nicely) weakly dominates x^{j+1} on $m^j(V^j)$, or y and x^{j+1} are payoff equivalent for i on $m^j(V^j)$.

Assume y (nicely) weakly dominates x^{j+1} on $m^j(V^j)$. Then, $(m^j)^{-1}(y)$ (nicely) weakly dominates $(m^j)^{-1}(x^{j+1})$ on V^j. So, define $V^{j+1}$by $V^j \backslash (m^j)^{-1}(x^{j+1})$, and let m^{j+1} be the restriction of m^j to V^{j+1}. Clearly (1)-(3) are satisfied for V^{j+1} and m^{j+1}.

Assume y and x^{j+1} are payoff equivalent on $m^j(V^j)$. Then, since y very weakly dominates r_i on $m^j(V^j)$, and since r_i nicely very weakly dominates x^{j+1} on $m^j(V^j)$, it must be that r_i and x^{j+1} are redundant on $m^j(V^j)$. Now, either $r_i \in m^j(V^j)$, or not.

Assume that $r_i \in m^j(V^j)$. Then, $(m^j)^{-1}(x^{j+1})$ and $(m^j)^{-1}(r_i)$ are redundant on V^j. So, let $V^{j+1} = V^j \backslash (m^j)^{-1}(x^{j+1})$, and let m^{j+1} be the restriction of m^j to V^{j+1}. Again, $(1) - (3)$ are clearly satisfied for V^{j+1} and m^{j+1}.

Assume that $r_i \notin m^j(V^j)$. Let $V^{j+1} = V^j$, and let m^{j+1} agree with m^j except that the strategy which used to map onto x^{j+1} is now mapped onto r_i. Once again, (1)-(3) are clearly satisfied for V^{j+1} and m^{j+1}, and we are done. ■

PROPOSITION 1. *Let X be a full reduction of S by nice weak dominance, and let Y be a full reduction of S by (nice) weak dominance. Then, after the removal of redundant strategies, Y is equivalent to a subset of X.*

Proof. X is a reduction of S by nice very weak dominance and Y is a reduction of S by (nice) very weak dominance. By Lemma C, a set $\hat{Y}$ equivalent to a subset of X is obtainable from Y by the iterative removal of strategies that are either (nicely) weakly dominated or redundant.

Since Y was a full reduction by (nice) weak dominance, $\hat{Y}$ differs from Y only by the removal of redundant strategies. So, after the removal of some redundant strategies, Y is equivalent to a subset of X. ■

The obvious application of Proposition 1 is the following:

THEOREM 1. *Let X and Y be full reductions of S by nice weak dominance. Then, X and Y are the same up to the addition or removal of redundant strategies and a renaming of strategies.*

Proof. Applying Proposition 1 twice, each of X and Y differs from a subset of the other by the removal of some redundant strategies and a renaming. ■

And, since weak dominance and nice weak dominance are equivalent for games that satisfy TDI, we have the following corollary:

COROLLARY 1. *Let (S,π) satisfy TDI, and let X and Y be full reductions of S*

by weak dominance. Then, X and Y are the same up to the addition or removal of redundant strategies and a renaming of strategies.

However, Proposition 1 has another nice implication. Even in games where TDI fails, Proposition 1 sets an upper bound on the amount by which different orders of elimination by weak dominance can matter. Even when TDI fails, the outcome of iterated nice weak dominance is (essentially) unique by Theorem 1. And by Proposition 1, any full reduction of S by weak dominance is equivalent to a subset of the outcome of iterated nice weak dominance. So, no matter what the order in which weak dominance is applied, the result will be at least as "tight" as the outcome of iterated nice weak dominance. In applications where weak dominance is important, but TDI fails, one should perhaps first apply nice weak dominance, since the result of this stage is independent of order, and only when all removals by nice weak dominance have been made consider (non-nice) weak dominance.[7]

V. Mixed Strategies

The results of the previous section dealt solely with eliminations by pure strategies. In this section we extend our analysis to mixed strategies. If W is a restriction of S, let $\sigma \in \Delta(W)$ indicate that σ is a mixed strategy profile in which each player i is using only pure strategies in W_i. In what follows, we will often write s_i where we more properly mean the mixed strategy that places probability 1 on s_i.

If a pure strategy s_i is weakly dominated on W, then every mixed strategy placing positive weight on s_i is also weakly dominated on W. So we initially consider only orders of elimination in which, when a pure strategy is removed, so is every mixed strategy using that pure strategy. At each stage, the set of remaining strategies is a subset of the pure strategies plus any mixtures over those pure strategies that have not been eliminated. For sets Ω of this form, a strategy s_i is weakly dominated on Ω if and only if it is weakly dominated versus the pure strategies underlying Ω.[8]

[7] Note that it is not the case that anything achievable by weak dominance can also be reached by first using nice weak dominance and then (not necessarily nice) weak dominance:

	L	R
T	1,1	2,0
B	1,3	1,1

Only R can be removed by NWD. Once R is gone, there are no further removals by WD. Conversely, under WD, one can remove both B and R simultaneously at the first step.

[8] To see this, note that s_i is weakly dominated on Ω if and only if there exists σ_i in the reconvexification of Ω that weakly dominates s_i on Ω. The mixture σ_i is in the reconvexification of Ω if and only if it is in the reconvexification of the set of pure strategies underlying Ω, and σ_i weakly dominates s_i on Ω if and only if σ_i weakly dominates s_i on the pure strategies underlying Ω.

Thus, for any sequence of sets of this form, one can consider a sequence in which at each stage a strategy is removed *if and only if* a pure strategy in its support is removed, except that when no more pure strategies are removable, one round of removing weakly dominated mixed strategies occurs. This new sequence will yield the same result. So, for sequences in which the removal of a pure strategy implies the removal of every mixed strategy using that pure strategy, the "order does not matter" result depends only on the order of removal of pure strategies. Thus it is this order to which we restrict attention in the analysis that follows. We later outline how the analysis can be extended to sequences in which pure strategies may be removed while leaving behind some mixed strategies using those pure strategies.

We begin by extending the definitions from Section IV to mixed strategies and by strengthening the TDI condition. We wish to construct a proof similar to that of Theorem 1. Recall that in that proof, we mapped each V_i^j into a subset of W_i^j. Here, we wish to do the same. However, because we are allowing eliminations of a pure strategy by a mixed strategy, it may happen at some stage that the strategy that some $s_i \in V_i^j$ was mapped into disappears because it is, for example, redundant to some mixture over W_i^j. To accommodate this, we must expand our notion of equivalent to a subset to allow members of V_i^j to be mapped onto mixed strategies in W_i^j. Recall also, that it was important in Theorem 1 that the images $m^j(V^j)$ could be thought of as coming from V by weak dominance, redundance or substitution. Since $m^j(V^j)$ itself may include mixed strategies, our definitions must thus include the possibility of a mixed strategy eliminating or replacing another mixed or pure strategy.

DEFINITION 7. For all $i \in N$, let V_i be a nonempty finite subset of $\Delta(S_i) \cup S_i$, and let $V = \bigcup_{i \in N} V_i$. Let $\sigma_i, \tau_i \in \Delta(S_i) \cup S_i$. Then

(i) σ_i *very weakly dominates** τ_i on V, written $\sigma_i\ VWD_V^*\ \tau_i$, if $\forall \gamma_{-i} \in \prod_{j \neq i} V_j \equiv V_{-i}$, $\pi_i(\sigma_i, \gamma_{-i}) \geq \pi_i(\tau_i, \gamma_{-i})$, and

(ii) σ_i *weakly dominates** τ_i on V, written $\sigma_i\ WD_V^*\ \tau_i$, if $\sigma_i\ VWD_V^*\ \tau_i$ and in addition for some $\gamma'_{-i} \in V_{-i}$, $\pi_i(\sigma_i, \gamma'_{-i}) > \pi_i(\tau_i, \gamma'_{-i})$.

As in Section IV, $\tau_i \in V_i$ is *very weakly dominated** on V if there is $\sigma_i \in \Delta(V_i \backslash \tau_i)$ where $\sigma_i\ VWD_V^*\ \tau_i$, and similarly for *weakly dominated**. Recall that V_i itself may contain mixed strategies. We abuse notation slightly by taking $\Delta(V_i)$ to be that subset of $\Delta(S_i)$ which can be implemented by randomizing over the elements of V_i. So, if $S_i = \{L, M, R\}$, and α_i is the mixed strategy given by $\alpha_i(L) = \alpha_i(M) = 1/2$, then $\Delta(\alpha_i, R)$ consists of all mixed strategies that place equal weight on L and M.

DEFINITION 8. A weak dominance* or very weak dominance* involving σ_i and τ_i on V is *nice* if for all $\gamma_{-i} \in V_{-i}$, $\pi_i(\sigma_i, \gamma_{-i}) = \pi_i(\tau_i, \gamma_{-i})$ implies $\pi(\sigma_i, \gamma_{-i}) = \pi(\tau_i, \gamma_{-i})$.

We define NWD_V^*, nicely weakly dominated* on V, and $NVWD_W^*$, nicely very weakly dominated* on W, in the obvious way.

OBSERVATION 4. For all $i \in N$, let W_i be a nonempty finite subset of $\Delta(S_i) \cup S_i$, and let $W = \bigcup_{i \in N} W_i$. Let $\sigma_i, \tau_i \in \Delta(S_i) \cup S_i$. For all $i \in N$, let V_i be a nonempty finite subset of $\Delta(W_i) \cup W_i$, and let $V = \bigcup_{i \in N} V_i$. Then $\sigma_i\ (N)VWD^*_W\ \tau_i$ implies $\sigma_i\ (N)VWD^*_V\ \tau_i$.

To see that niceness is inherited, note that by the definition of nice (very) weak dominance* on W, $\pi_i(\sigma_i, y_{-i}) \geq \pi_i(\tau_i, y_{-i})$ for all $y_{-i} \in \prod_{k \in N \setminus i} W_k$. So, if $\pi_i(\sigma_i, \omega_{-i}) = \pi_i(\tau_i, \omega_{-i})$ for some $\omega_{-i} \in \prod_{k \in N \setminus i} \Delta(W_k)$, then $\pi_i(\sigma_i, y_{-i}) = \pi_i(\tau_i, y_{-i})$ for all y_{-i} in the support of ω_{-i}. But, then because of niceness, $\pi(\sigma_i, y_{-i}) = \pi(\tau_i, y_{-i})$ for all y_{-i} in the support of ω_{-i}, and so $\pi(\sigma_i, \omega_{-i}) = \pi(\tau_i, \omega_{-i})$.

For a restriction W of S, a reduction of W by (nice) (very) weak dominance* is defined as in the previous section to be a set that can be obtained from W by iteratively eliminating (nicely) (very) weakly dominated* pure strategies. Similarly, a one-at-a-time reduction of W by (nice) (very) weak dominance* is a reduction of W by (nice) (very) weak dominance* in which only one pure strategy is eliminated at each round.

DEFINITION 9. For all $i \in N$, let V_i be a nonempty finite subset of $\Delta(S_i) \cup S_i$, and let $V = \bigcup_{i \in N} V_i$. Let $\sigma_i, \tau_i \in \Delta(S_i) \cup S_i$. Then σ_i is *redundant* to* τ_i on V, written $\sigma_i\ R^*_V\ \tau_i$ if for all $\gamma_{-i} \in V_{-i}$, $\pi(\sigma_i, \gamma_{-i}) = \pi(\tau_i, \gamma_{-i})$. A strategy τ_i is *redundant* on* V if there is $\sigma_i \in \Delta(V_i \setminus \tau_i)$ with $\sigma_i\ R^*_V\ \tau_i$.

DEFINITION 10. (S, π) satisfies *TDI** if for all restrictions W and for all s_i, if s_i is very weakly dominated* on W, then it is either weakly dominated* on W or redundant* on W.

We were able to state TDI in terms of a simple condition on the payoffs to pure strategies in a game. The same condition on payoffs is not enough to imply TDI*. Consider G3.

		II L	II C	II R
	T	2,1	2,3	0,2
I	M	0,3	3,1	0,2
	M′	1,4	1,4	1,4
	B	1,4	0,3	0,2

Fig. G3.

G3 satisfies TDI, but TDI* fails because R is very weakly dominated* on $S \setminus B$ by $\frac{1}{2}L + \frac{1}{2}C$, but R is neither weakly dominated* nor redundant* on $S \setminus B$. Furthermore, a strengthening of TDI is clearly necessary because, as this game illustrates, even when

TDI holds, the order of elimination under WD^* can matter: if one first removes B, then $S\backslash B$ is a full reduction of S by weak dominance*, while if one first removes R, then B and M' can also be removed. But, $S\backslash B$ and $\{T, M\} \times \{L, C\}$ are clearly not equivalent.

G3 has a highly non-generic feature: R is weakly dominated* by $\frac{1}{2}L + \frac{1}{2}C$, but by no other mixture. Consider different assignments of payoffs, subject to existing ties in payoffs for each player being maintained (each player receives four different payoffs, so these assignments correspond to elements of $\Re^8$). For all such assignments, the equality of payoffs in the third row will remain, and so in particular, any mixture of L and C will give the same payoff as R versus M'. However, for almost all such assignments, one of two things will occur: either there will be a mixture of L and C that strictly dominates* R on $\{T, M, B\}$, or there will be no mixture that even weakly dominates* R on $\{T, M, B\}$. In particular, the set of payoff assignments yielding a mixture of L and C that weakly dominates* R on $\{T, M, B\}$ but not yielding any mixture that strictly dominates* R on $\{T, M, B\}$ is a lower dimensional subspace of $\Re^8$. If one allows perturbations that do not respect some of the existing ties, then the situation in this game becomes even less likely.

So, even allowing for some "structural" ties in payoffs, for almost all games, if there is $\sigma_i \in \Delta(W_i\backslash s_i)$ with $\sigma_i\ VWD^*_W\ s_i$, then there is another mixed strategy $\tau_i \in \Delta(W_i\backslash s_i)$ that strictly dominates* s_i except versus those opposition strategy profiles s_{-i} on which all elements in the support of τ_i give the same payoff to i as does s_i. For such strategy profiles, TDI is enough to imply TDI*. So, if a game satisfies TDI, then it will generically satisfy TDI* as well. Thus, in particular, TDI* will be generically satisfied for the normal form of any given extensive form and for the discrete first price auction. See Appendix B for a formal statement and proof of this result.

For games satisfying TDI*, weak dominance* is equivalent to nice weak dominance*, so the analysis of Section IV goes through fairly directly:

LEMMA A*. *Let W be a reduction of V by (nice) very weak dominance*. Then W is a one-at-a-time reduction of V by (nice) very weak dominance*.*

Proof. As in the proof of Lemma A. ■

OBSERVATION 5. $\sigma_i\ NVWD^*_W\ s_i$ if and only if either $\sigma_i\ NWD^*_W\ s_i$ or $\sigma_i\ R^*_W\ s_i$.

DEFINITION 11. Let W and V be restrictions of S.[9] V is *equivalent* to a subset of* W if there exists a one-to-one map m such that $\forall i \in N$ and $\forall s_i \in V_i$, $m(s_i) \in \Delta(W_i) \cup W_i$ and $\forall s \in \prod_{i\in N} V_i$, $\pi(s) = \pi(m(s_1), ..., m(s_n))$.

[9]Recall that by the definition of restriction, W and V contain only pure strategies.

DEFINITION 12. Let W be a restriction of S, and $\forall i \in N$, let V_i be a non-empty subset of $\Delta(S_i) \cup S_i$. Let $V \equiv \bigcup_{i \in N} V_i$. Let $V^0 \equiv W$. Then, the set V is a *reduction of W by (nice) weak dominance*/redundance*/substitution** if V can be obtained from V^0 by performing a finite number of iterations of the following process: obtain V^{j+1} from V^j by for some $i \in N$ either (i) letting $V^{j+1} = V^j \backslash \tau_i^j$, where τ_i^j is (nicely) weakly dominated* on V^j by an element of $\Delta(V_i^j \backslash \tau_i^j)$, or (ii) letting $V^{j+1} = V^j \backslash \tau_i^j$, where τ_i^j is redundant* on V^j to an element of $\Delta(V_i^j \backslash \tau_i^j)$, or (iii) letting $V^{j+1} = (V^j \backslash \tau_i^j) \cup \sigma_i$, where $\sigma_i \in \Delta(W_i) \backslash V_i^j$ and σ_i is redundant* on V^j to τ_i^j.

Note that any set V obtained from a restriction W of S by (nice) weak dominance*/ redundance*/substitution* is a finite set of pure or mixed strategies with at least one strategy for each player, and that only the substitution* operation creates new strategies.

LEMMA B*. *Let W be a restriction of S, let T be a reduction of W by (nice) weak dominance*/redundance*/substitution*, and let $\gamma \in \Delta(W_i) \cup W_i$ for some i. Then there exists $\gamma' \in \Delta(T_i)$ such that γ' (nicely) very weakly dominates* γ on T.*

Proof. Let $\gamma \in \Delta(W_i) \cup W_i$. Since T is a reduction of W by (nice) weak dominance*/redundance*/substitution*, there exist sets $T^0, T^1, \ldots, T^m$ such that $T^0 = W$ and $T^m = T$ and each T^{j+1} is obtained from T^j as in the definition of a reduction by (nice) weak dominance*/redundance*/substitution*. The Lemma is satisfied for T^0 because $\gamma \in \Delta(T_i^0)$, so γ itself serves (if $\gamma \in S_i$, then the mixed strategy which places weight 1 on γ_i serves). Assume that the Lemma holds for T^j. Then there exists $\gamma' \in \Delta(T_i^j)$ such that γ' (nicely) very weakly dominates* γ on T^j. If $\gamma' \in \Delta(T^{j+1})$, then γ' serves and the Lemma holds for T^{j+1}.

Suppose that $\gamma' \notin \Delta(T_i^{j+1})$. Then γ' places positive weight on the strategy τ_i^j that was eliminated when T^{j+1} was obtained from T^j. If τ_i^j was eliminated by (nice) weak dominance* or redundance* (no strategy added) by some $\zeta \in \Delta(T_i^j \backslash \tau_i^j)$, then let γ'' be the strategy that places weight $\gamma'(\sigma) + \zeta(\sigma)\gamma'(\tau_i^j)$ on strategies $\sigma \in T_i^j \backslash \tau_i^j$. If τ_i^j was eliminated and replaced by redundant* strategy $\zeta' \in \Delta(W_i \backslash \tau_i^j)$, then let γ'' be the strategy that places weight $\gamma'(\sigma)$ on strategies $\sigma \in T_i^j \backslash \tau_i^j$ and weight $\gamma'(\tau_i^j)$ on ζ'. Then $\gamma'' \in \Delta(T_i^{j+1})$ and γ'' (nicely) very weakly dominates γ on T^{j+1}. The result follows by induction. ■

Let $x_i \in S_i$, and let $\sigma_i, \rho_i \in \Delta(S_i)$ be such that $\rho_i(x_i) = 0$. Then, define $\sigma_i\{x_i \to \rho_i\}$ as the mixed strategy γ_i given by $\gamma_i(s_i) = \sigma_i(s_i) + \sigma_i(x_i)\rho_i(s_i)$ for $s_i \neq x_i$, and $\gamma_i(x_i) = 0$. That is, $\sigma_i\{x_i \to \rho_i\}$ is obtained from σ_i by redistributing the weight on x_i according to ρ_i.

LEMMA C*. *Let W be a restriction of S, let $\hat{W}$ be a reduction of W by nice very weak dominance*, and let V be a reduction of W by (nice) very weak dominance*.*

Then there exists $\hat{V}$ equivalent to a subset of $\hat{W}$, where $\hat{V}$ is obtainable from V by the iterative removal of strategies that are either (nicely) weakly dominated* or redundant*.*

Proof. The proof is similar to that of Lemma C. There are two main differences. First, in Lemma C, each V_i^k was mapped onto a subset of W_i^k. In this case, each V_i^k is mapped onto a subset of mixtures over W_i^k. Second, because of the mixed strategies, the step at which we show that x^{k+1} can be eliminated is more intricate.

Since $\hat{W}$ is a reduction of W by nice very weak dominance*, by Lemma A* we can write $\hat{W} = W\backslash x^1, ..., x^m$, where for $k = 1, ..., m$, $W^k = W\backslash x^1, ..., x^k$ is a (one-at-a-time) reduction of $W\backslash x^1, ..., x^{k-1}$ by nice very weak dominance*. Let $W^0 = W$.

We proceed by induction. For some $j \in \{0, ..., k-1\}$ assume that we have V^j and m^j such that

(1) $V^j \subseteq V$ can be obtained from V by iteratively eliminating (pure) strategies that are either (nicely) weakly dominated* or redundant*.

(2) m^j is a one to one map such that $\forall i \in N$, $\forall s_i \in V_i^j$, $m^j(s_i) \in \Delta(W_i^j) \cup W_i^j$ and $\forall t \in \prod_{i \in N} V_i^j$, $\pi(t) = \pi(m^j(t_1), ..., m^j(t_n))$, and

(3) $m^j(V^j)$ can be obtained iteratively from V by eliminating strategies that are either (nicely) weakly dominated* or redundant* or by substitution*.

So, while V^j remains a subset of pure strategies, its image in W^j may involve mixtures. Given V^j and m^j satisfying $(1)-(3)$, we will exhibit V^{j+1} and m^{j+1} satisfying $(1)-(3)$. This will thus establish the Lemma, since conditions $(1)-(3)$ are satisfied for $j = 0$ by taking $V^0 = V$ and m^0 as the identity map.

So, given V^j and m^j, let us construct V^{j+1} and m^{j+1}. Let i be the player to whom x^{j+1} belongs and let $\rho_i \in \Delta(W_i^{j+1})$ be a strategy that nicely very weakly dominates* x^{j+1} on W^j. Then, by Observation 4, ρ_i nicely very weakly dominates* x^{j+1} on $m^j(V^j)$.

So, by Observation 5, either ρ_i nicely weakly dominates* x^{j+1} on $m^j(V^j)$ or ρ_i and x^{j+1} are redundant* on $m^j(V^j)$. Assume ρ_i and $x^{j+1}$are redundant* on $m^j(V^j)$. Then, for each $s_i \in V_i^j$, $m^j(s_i)\{x^{j+1} \to \rho_i\}$ is redundant* to $m^j(s_i)$ on $m^j(V^j)$. Now, it may turn out that there are sets of strategies in V_i^j all elements of which are mapped into the same strategy when x^{j+1} is replaced by ρ_i. If so, then these strategies are redundant* on V^j. Construct V_i^{j+1} by removing all but one of any such set, and let $V_k^{j+1} = V_k^j$ for $k \neq i$. For $s_i \in V^{j+1}$, define $m^{j+1}(s_i) = m^j(s_i)\{x^{j+1} \to \rho_i\}$, while for $k \neq i$, and $s_k \in V_k^j$, let $m^{j+1}(s_k) = m^j(s_k)$. Then, $(1)-(3)$ are clearly satisfied. In particular, $m^{j+1}(V^{j+1})$ is obtained from $m^j(V^j)$ by substitution* and (if more than one strategy in V_i^j is mapped onto the same strategy when x^{j+1} is replaced by ρ_i) by redundance*.

Assume next that ρ_i nicely weakly dominates* x^{j+1} on $m^j(V^j)$. Let $\sigma_i \in m^j(V_i^j)$ have $\sigma_i(x^{j+1}) > 0$. We claim first that there is $\gamma_i \in \Delta(m^j(V_i^j))$ such that $\gamma_i(x^{j+1}) = 0$, and such that γ_i (nicely) weakly dominates σ_i. To see this, note first that σ_i is nicely weakly dominated* by $\eta_i \equiv \sigma_i\{x^{j+1} \to \rho_i\}$. By Lemma B*, there is $\phi_i \in \Delta(m^j(V_i^j))$

which (nicely) very weakly dominates* η_i and therefore (nicely) weakly dominates* σ_i. Let Θ be the set of elements of $\Delta(m^j(V_i^j))$ which (nicely) very weakly dominate* η_i. Θ is non-empty since it contains η_i. It is also clearly compact. So, there is $\gamma_i \in \arg\max_{\alpha_i \in \Theta} f(\alpha_i)$, where

$$f(\alpha_i) \equiv \sum_{\delta_{-i} \in \prod_{k \neq i} m^j(V_k^j)} \pi_i(\alpha_i, \delta_{-i}).$$

Clearly γ_i (nicely) weakly dominates* σ_i, and is in $\Delta(m^j(V_i^j))$. Assume $\gamma_i(x^{j+1}) > 0$. Then, since ρ_i nicely weakly dominates* x^{j+1} on $m^j(V^j)$, $\gamma_i\{x^{j+1} \rightarrow \rho_i\}$ (nicely) weakly dominates* γ_i on $m^j(V^j)$. And, by Lemma B*, there is $\mu_i \in \Delta(m^j(V_i^j))$ which (nicely) very weakly dominates* $\gamma_i\{x^{j+1} \rightarrow \rho_i\}$. But, then $\mu_i \in \Theta$, and $f(\mu_i) > f(\gamma_i)$, a contradiction.

Let $Y_i \equiv \{s_i \in V_i^j | m^j(s_i)(x^{j+1}) > 0)$. By the previous claim, for each $s_i \in Y_i$, there is $\gamma_i \in \Delta(m^j(V^j))$ such that γ_i (nicely) weakly dominates* $m^j(s_i)$, and such that $\gamma_i(x^{j+1}) = 0$. But, then it must be the case that $\gamma_i(m^j(Y_i)) = 0$, and so $(m^j)^{-1}(\gamma_i) \in \Delta(V_i \backslash Y_i)$. So, define $V_i^{j+1} = V_i^j \backslash Y_i$, $V_k^{j+1} = V_k^j$ for $k \neq i$, and define m^{j+1} as the restriction of m^j to V^{j+1}. By the preceding argument, each s_i removed from V_i^j was (nicely) weakly dominated* by a mixture over V_i^{j+1}, and each σ_i removed from $m^j(V_i^j)$ is (nicely) weakly dominated* by a mixture over $m^{j+1}(V_i^{j+1})$, and so (1)-(3) again clearly hold. ■

PROPOSITION 2. *Let X be a full reduction of S by nice weak dominance*, and let Y be a full reduction of S by (nice) weak dominance*. Then, after the removal of redundant* strategies, Y is equivalent to a subset of X.*

Proof of Proposition. As in the proof of Proposition 1. ■

As before, we have:

THEOREM 2. *Let X and Y be full reductions of S by nice weak dominance*. Then, X and Y are the same up to the addition or removal of redundant* strategies and a renaming of strategies.*

Proof. As for Theorem 1. ■

COROLLARY 2. *Let (S,π) satisfy TDI*, and let X and Y be full reductions of S by weak dominance*. Then, X and Y are the same up to the addition or removal of redundant* strategies and a renaming of strategies.*

We now sketch how the analysis of this section can be extended to allow nicely weakly dominated* pure strategies to be eliminated without necessarily eliminating all mixtures that use those pure strategies. Let $\Omega^0, \Omega^1, \ldots, \Omega^m$, where $\Omega^0 = \Delta(S)$, be a finite sequence of rounds of elimination of nicely weakly dominated* mixed strategies such that Ω^m is a full reduction of S by nice weak dominance*.[10] For each k, let W^k be the set of pure strategies s for which there is a mixed strategy in Ω^k that puts probability 1 on s. Then, a fairly straightforward induction argument establishes that any strategy in Ω^k that uses a pure strategy not in W^k is nicely very weakly dominated* on Ω^k by some mixed strategy using only elements of W^k. Consider the sequence $\{\Psi^k\}$ obtained by removing from each Ω^k any mixed strategy using a pure strategy not in W^k. By the previous assertion, this is a valid sequence of removals by nice very weak dominance*. By Observation 5, any removal that was not by nice weak dominance* was redundant*, and so the proof of Proposition 2 establishes that Ψ^m, if it contains no nicely weakly dominated* strategies, depends on the order of removals only up to redundancy* and renaming. So, we must show that Ψ^m contains no nicely weakly dominated* strategies. Consider any strategy $\sigma \in \Omega^m \backslash \Psi^m$. By the assertion, σ is nicely very weakly dominated* by some mixture over W^m. Since Ω^m is a full reduction of S by nice weak dominance*, σ is not nicely weakly dominated*, and so by Observation 5, σ is redundant* to some mixture over W^m. But then, a strategy in $\Omega^m$is nicely weakly dominated* on Ω^m if and only if it is nicely weakly dominated* on Ψ^m, which, since Ω^m is a full reduction of S by nice weak dominance*, establishes that Ψ^m is also a full reduction of S by nice weak dominance*, establishing the result.

VI. Weak Dominance and Backward Induction

In its purest form, backward induction consists of iteratively removing actions that are strictly dominated given the information available when that action is taken.

A normal form strategy that is consistent with reaching a particular information set and takes a strictly dominated action at that information set is weakly dominated by one that differs only by taking a dominating action at that information set. So, any sequence of removals of actions by backward induction in the extensive form corresponds to a sequence of sets of removals of strategies by weak dominance in the normal form.[11]

This relationship between backward induction in the extensive form and weak dominance in the normal form extends to their respective motivations. Backward induction requires that the action chosen at an information set must not be strictly dominated when the opponents play in such a way that the information set is reached. Weak dominance requires that a chosen strategy not be strictly dominated by another strategy when the opponents play in such a way that the choice between these two

[10] Dealing with infinite sequences of removals introduces issues which we would prefer not to deal with here, although the main ideas of this analysis translate directly.

[11] Glazer and Rubinstein (1993) start from this observation in defining "guides," which are essentially prespecified orders in which to check weak dominance.

strategies affects the decision maker's payoffs. So, each is about making decisions under the (possibly counterfactual) hypothesis that the opponents have not chosen in a way such that the decision "does not matter."[12]

Thus, at an intuitive level, there seems to be an intimate relationship between backward induction and weak dominance. However, the order of removals under weak dominance can matter, while backward induction is deterministic.[13] This raises the possibility that backward induction and weak dominance differ at some fundamental level that we have failed to understand. What is it about these two concepts that leads one to be deterministic and the other not?

We claim that the distinction between weak dominance and backward induction comes down to a fairly minor difference in what the two concepts mean by "does not matter." For backward induction, "does not matter" means "makes the information set unreachable," while for weak dominance, "does not matter" means "makes the choice irrelevant for the decision maker's payoffs." Of course, if the information set is not reached, the choice is irrelevant for the decision maker's payoffs *and* every other player's payoffs. Thus, in backward induction a strategy profile for players other than i is excluded in player i's decision making process only if, regardless of the decision i makes, the profile gives the same payoffs to all players. So, eliminations in backward induction are eliminations by nice weak dominance and thus are order independent.

Thus, the fact that backward induction is deterministic while weak dominance is not does not reflect some fundamental difference in their motivations, but rather the fairly simple difference between nice weak dominance and weak dominance.

The connection between backward induction and nice weak dominance allows us to strengthen some of the results of Rochet (1980). Rochet shows that for extensive form games with perfect information satisfying an extensive form version of (1), the unique backward induction payoff is the same as the unique payoff from iterated weak dominance in the normal form, which is the same as the unique proper equilibrium payoff. This result extends to games satisfying an extensive form version of TDI. To see this, note that an extensive form game with perfect information that satisfies TDI has a unique backward induction payoff. Since backward induction eliminations correspond to nice weak dominance eliminations, by Proposition 1, any full reduction by nice weak dominance or weak dominance yields the unique backward induction payoff. Since any proper equilibrium is a backward induction solution, any proper equilibrium also has the same payoff.

In addition, Proposition 1 shows that if an extensive form game has a unique backward induction payoff, then any full reduction by nice weak dominance or weak

[12] On the general relation between normal and extensive form motivations and the implementations of solution ideas, including a discussion of counterfactual reasoning in normal form games, see Mailath *et al.* (1993).

[13] The set of outcomes that survive backward induction remains deterministic even if we allow iterative elimination in any order of dominated actions at information sets, i.e. even if the elimination process does not start at the end of the tree and move up.

dominance also contains only that payoff, and any violations of TDI are non-essential in that they do not affect the outcome of iterated weak dominance.

VII. Weak Dominance and Complexity

We close with a brief comment on how our results interact with those of Gilboa *et al.* (1993). They point out that, in general, computational problems involving weak dominance are hard. In particular, given a full reduction of a game, the question remains whether different choices earlier in the sequence of weak dominance removals might have led to a strategically different result. To figure out all the strategic implications of weak dominance, one must thus check all possible orders of removal, which cannot be done in polynomial time. Gilboa *et al.* interpret this result as casting additional doubt on the use of weak dominance as a solution idea.[14]

Consider, however, nice weak dominance in general games or weak dominance in games that satisfy TDI or some other condition such that order does not matter. Then, once one has arrived at a full reduction, one knows no other order could have resulted in a strategically different game. Since finding a full reduction is a polynomial problem, nice weak dominance avoids some of the computational problems of weak dominance and so for an important class of games weak dominance becomes less suspect.[15]

Appendix A

In this section we formalize the argument that the discrete first price auction generically satisfies TDI. Think of the auction as being generated by first fixing a set of players $1, \ldots, n$, a finite set B of admissible bids, a set of names of signals $\Omega = \{\omega^1, \ldots, \omega^m\}$, a measure ρ on Ω^n, and maps $V_i : \Omega^{n-1} \times \Omega \to \Re$ giving the value of the object to player i when the other bidders receive signals $\omega_{-1} \in \Omega^{n-1}$ and i receives signal $\omega_i \in \Omega$ (notation aside, having different sets Ω for different players presents no difficulties). Since V_i assigns a value to each of m^n different signal profiles,

[14] See Samuelson (1992) on another concern with the concept of iterated weak dominance, namely that it cannot be grounded by assuming that it is common knowledge that players do not play weakly dominated strategies.

[15] See Gilboa *et al.* (1993) for a formal development of this material. In particular, see Section 4.3 for a proof that a full reduction can be computed in polynomial time. The reader may also wonder whether checking TDI is a polynomial problem. The answer is yes, because checking where the antecedent of the TDI condition holds involves checking, for each player i, $\frac{1}{2} \mid S_i \mid\mid S_i - 1 \mid\mid S_{-i} \mid$ possible equalities (this is the number of r_i, s_i, s_{-i} triplets) and then for each such equality, checking $n-1$ further equalities, which is clearly a polynomial problem.

the function V_i can be associated in the obvious way with an element of $\Re^{m^n}$. Each player's strategy space is the set of all maps from signals in Ω to bids in B, and so is finite. Fix a pure strategy profile s_{-i} for the players other than i. Given this behavior of the other players and any pure strategy s_i for i, let $p_i(\omega, s(\omega))$ be i's probability of winning when signals are ω and bids are according to s. The expected payoff to i from following s_i is thus

$$\pi_i(s_i, s_{-i}) = \sum_{\omega \in \Omega^n} \rho(\omega) p_i(\omega, s(\omega)) \left(V_i(\omega) - s_i(\omega_i)\right).$$

This expression depends on V_i only directly (ρ and ρ_i do not depend on the value assignment), and so it is a linear function of V_i. Consider two pure strategies s_i and t_i for i and a pure strategy profile s_{-i} for i's opponents. Suppose that $\pi_i(s_i, s_{-i}) = \pi_i(t_i, s_{-i})$.

Case (i): $\exists \omega$ such that $\rho(\omega) > 0$ and $p_i(\omega, s(\omega)) \neq p_i(\omega, t_i(\omega_i), s_{-i}(w_{-i}))$. Then the equation $\pi_i(s_i, s_{-i}) - \pi_i(t_i, s_{-i}) = 0$ has nonzero coefficient on $V_i(\omega)$ and so is satisfied for a set of possible value assignments V_i which is a lower-dimensional subspace of $\Re^{m^n}$.

Case (ii): $\forall \omega$ such that $\rho(\omega) > 0$, $p_i(\omega, s(\omega)) = p_i(\omega, t_i(\omega_i), s_{-i}(\omega_{-i}))$. If $\rho(\omega) > 0$ implies $s_i(\omega_i) = t_i(\omega_i)$, then clearly $\pi(s_i, s_{-i}) = \pi(t_i, s_{-i})$. Suppose $\rho(\hat{\omega}) > 0$ and $s_i(\hat{\omega}_i) \neq t_i(\hat{\omega}_i)$ for some $\hat{\omega}$. W.l.o.g., assume $s_i(\hat{\omega}_i) > t_i(\hat{\omega}_i)$. By assumption, i's probability of winning facing $s_{-i}(\hat{\omega}_{-i})$ is the same with $s_i(\hat{\omega}_i)$ and $t_i(\hat{\omega}_i)$. But, then the highest bid for the opponents under $\hat{\omega}_{-i}$ must be either less than $t_i(\hat{\omega}_i)$ or greater than $s_i(\hat{\omega}_i)$. In either case, player i's change from $s_i(\hat{\omega}_i)$ to $t_i(\hat{\omega}_i)$ does not affect the payoff of players other than i.

Note that there are a finite number of s_i, t_i, s_{-i} combinations, and so the set of V_i for which the situation of Case (i) holds is a finite union of lower-dimensional subspaces of $\Re^{m^n}$. In the situation of Case (ii), TDI holds.

Appendix B

In this appendix, we formalize the notion that TDI generically implies TDI*. We first need to formalize the idea of generic payoffs subject to any given set of payoff ties. To do this, retain the previous notation for players and strategies, but add for each player i, a finite set O_i and two functions, $\phi_i : S \to O_i$ and $\rho_i : O_i \to \Re$. In games as we have considered them so far, π_i is the composition of ϕ_i with ρ_i. Note that any given ρ_i can be thought of as a member of $\Re^{O_i}$. Genericity statements can then be made at the level of ρ_i. We interpret O_i as a set of descriptions of economic outcomes player i may receive in the game. ϕ_i describes which outcome is associated with each strategy profile in the game, and finally ρ_i describes i's utility over outcomes. In this interpretation, if $\phi_i(s) = \phi_i(t)$, then from player i's point of view, s and t lead to identical outcomes, and so whatever his utility function over outcomes, he will be indifferent over s and t. On the other hand, if $\phi_i(s) \neq \phi_i(t)$, then the outcomes associated with s and t are distinct, and so i is indifferent over s and t only if his utility function happens to be indifferent over $\phi_i(s)$ and $\phi_i(t)$. For example, in a private value auction, the elements of O_i might describe whether i won the object or not and how much was paid. ρ_i would then capture information about i's attitudes toward money and the object. For any such utility function, i is indifferent over strategy profiles in which he doesn't win, but only for very specific realizations of price will i be indifferent between losing and winning.

These preliminaries out of the way, we can state and prove our main result:

THEOREM. *Fix N, and for each $i \in N$, fix S_i, and ϕ_i. Assume that TDI is satisfied at the level of outcomes. That is, assume that for all $i, j \in N$, $r_i, t_i \in S_i$, and $s_{-i} \in S_{-i}$,*

$$\phi_i(r_i, s_{-i}) = \phi_i(t_i, s_{-i}) \Rightarrow \phi_j(r_i, s_{-i}) = \phi_j(t_i, s_{-i}).$$

Let F be the set of ρ such that TDI fails. Then, $Cl(F)$ has empty interior, where $Cl(.)$ indicates closure.*

Proof. Let W be a restriction of S, and let $t_i \in S_i \backslash W_i$ for some i. Let $C(.)$ indicate the carrier of a mixed strategy. We will show that the set of ρ_i such that t_i is VWD^*_W by some σ_i with $C(\sigma_i) = W_i$ but t_i is neither R^*_W nor WD^*_W by any γ_i with $C(\gamma_i) = W_i$ has closure with empty interior. Since there are a finite number of (W, t_i) pairs, this establishes the result.

Assume that $\exists s_{-i} \in W_{-i}$ such that $\forall s_i \in W_i$, $\phi_i(s_i, s_{-i}) = \phi_i(t_i, s_{-i})$. Then, by TDI, ϕ_j will also be constant, and so regardless of σ_i (with $C(\sigma_i) = W_i$), and regardless of ρ, $\pi(\sigma_i, s_{-i}) = \pi(t_i, s_{-i})$ will hold. So, no failure of TDI* relative to t_i and W could ever be traced to s_{-i}. So, without loss of generality, assume that W has been chosen in such a way that $\nexists\, s_{-i} \in W_{-i}$ such that $\forall s_i \in W_i$, $\phi_i(s_i, s_{-i}) = \phi_i(t_i, s_{-i})$.

Define A_{no} as the set of all ρ_i such that t_i is not VWD^* on W by any σ_i with $C(\sigma_i) = W_i$. Define A_{strict} as the set of ρ_i such that t_i is strictly dominated on W by some σ_i with $C(\sigma_i) = W_i$.[16] Note that A_{strict} is open, since if σ_i strictly dominates t_i on W for ρ_i, then it continues to do so for all ρ_i' in a neighborhood of ρ_i. Let $A_{no} \cup A_{strict} \equiv A$. Then, any failure of TDI* is contained in the set A^c. So, it is enough to show that $Cl(A^c)$ contains no open sets.

Let $\rho_i \in A^c$. We will examine two cases.

Case (i): Assume that for some $Y_{-i} \subseteq W_{-i}$, $\phi_i(W_i \times Y_{-i}) \subseteq \phi_i(t_i \times Y_{-i})$. Then, let

$$a \in \arg\max_{a' \in \phi_i(t_i \times Y_{-i})} \rho_i(a'),$$

and let $r_{-i} \in Y_{-i}$ be such that $\phi_i(t_i, r_{-i}) = a$. Generate ρ_i' from ρ_i by increasing $\rho_i(a)$ by any small positive amount, and otherwise leaving ρ_i unchanged. Then,

[16] A strategy $\sigma_i \in \Delta(S_i)$ strictly dominates $s_i \in S_i$ on W if for all $s_{-i} \in W_{-i}$, $\pi_i(\sigma_i, s_{-i}) > \pi_i(s_i, s_{-i})$.

$\forall s_i \in W_i$, $\rho_i'(\phi_i(s_i, r_{-i})) \leq \rho_i'(\phi_i(t_i, r_{-i}))$ with strict inequality at least once, since for at least some $s_i \in W_i$, $\phi_i(s_i, r_{-i}) \neq a$. But, then for any σ_i having $C(\sigma_i) = W_i$, $\pi_i'(\sigma_i, r_{-i}) < \pi_i'(t_i, r_{-i})$, where π' is the payoff function for i generated by ϕ_i and ρ_i'. So, for ρ_i', t_i is not VWD^* on W by any σ_i with $C(\sigma_i) = W_i$, and so $\rho_i' \in A_{no}$. Further, this remains true of any perturbation ρ_i'' of ρ_i' such that $a = \arg\max_{a' \in \phi_i(t_i \times Y_{-i})} \rho_i''(a')$, and so for all ρ_i'' in a neighborhood of ρ_i', $\rho_i'' \in A_{no}$. That is, ρ_i' is in $Int(A_{no})$, where $Int(.)$ indicates interior. Thus, any neighborhood of ρ_i has non-empty intersection with $Int(A_{no})$. Since $Cl(A^c) \cap Int(A_{no}) = \emptyset$, this implies that there is no neighborhood of ρ_i contained in $Cl(A^c)$.

Case (ii): Assume that for each $Y_{-i} \subseteq W_{-i}$, $\phi_i(W_i \times Y_{-i}) \nsubseteq \phi_i(t_i \times Y_{-i})$. Let σ_i be such that $C(\sigma_i) = W_i$ and such that $\forall s_{-i} \in W_{-i}$, $\pi_i(\sigma_i, s_{-i}) \geq \pi_i(t_i, s_{-i})$, where π_i is the payoff function generated by ρ_i. Pick some $a \in \phi_i(W_i \times W_{-i}) \backslash \phi_i(t_i \times W_{-i})$. Generate ρ_i' from ρ_i by increasing $\rho_i(a)$ by any small positive amount, and otherwise leaving ρ_i unchanged. Then, $\pi_i'(\sigma_i, s_{-i}) > \pi_i'(t_i, s_{-i})$ for any $s_{-i} \in W_{-i}$ such that $a \in \phi_i(W_i \times s_{-i})$, of which there is at least one, while $\forall s_{-i} \in W_{-i}$, $\pi_i'(\sigma_i, s_{-i}) \geq \pi_i'(t_i, s_{-i})$. Let W'_{-i} be the subset of W_{-i} such that $\forall s_{-i} \in W'_{-i}$, $\pi_i'(\sigma_i, s_{-i}) = \pi_i'(t_i, s_{-i})$. If W'_{-i} is empty, then we are done. Otherwise, choose $b \in \phi_i(W_i \times W'_{-i}) \backslash \phi_i(t_i \times W'_{-i})$. Generate ρ_i'' from ρ_i' by increasing $\rho_i'(b)$ by a small positive amount, and otherwise leaving ρ_i' unchanged. Then, $\pi_i''(\sigma_i, s_{-i}) > \pi_i''(t_i, s_{-i})$, for any $s_{-i} \in W'_{-i}$ such that $b \in \phi_i(W_i \times s_{-i})$, of which there is at least one, while $\forall s_{-i} \in W'_{-i}$, $\pi_i''(\sigma_i, s_{-i}) \geq \pi_i''(t_i, s_{-i})$. And, by making $\rho_i''(b) - \rho_i'(b)$ small enough, the strict inequalities for s_{-i} in $W_{-i} \backslash W'_{-i}$ will be retained. Proceeding in this way, we generate $\rho_i^{\prime\cdots\prime}$ arbitrarily close to ρ_i under which σ_i strictly dominates t_i on W_{-i}, and so $\rho_i^{\prime\cdots\prime} \in A_{strict}$. Since A_{strict} is open, we have again shown that there is no neighborhood of ρ_i contained in $Cl(A^c)$. ■

References

Gilboa, I., Kalai, E., and Zemel, E. (1990). "On the Order of Eliminating Dominated Strategies," *Oper. Res. Lett.* **9**, 85-89.

Gilboa, I., Kalai, E., and Zemel, E. (1993). "The Complexity of Eliminating Dominated Strategies," *Math. Oper. Res.* 553-565.

Glazer, J. and Rubinstein, A. (1993). "Simplicity of Solution Concepts: Subgame Perfect Equilibrium in Extensive Games vs. Iteratively Undominated Strategies in Normal Games," Sackler Institute of Economic Studies Working Paper, 9-93.

Gretlein, R. J. (1983). "Dominance Elimination Procedures on Finite Alternative Games," *Int. J. Game Theory* **12**, 107-113.

Mailath, G., Samuelson, L., and Swinkels, J. (1993). "Extensive Form Reasoning in Normal Form Games," *Econometrica* **61**, 273-302.

Moulin, H. (1979). "Dominance Solvable Voting Schemes." *Econometrica* **47**, 1337-1351.

Moulin, H. (1984). "Dominance Solvability and Cournot Stability." *Math. Soc. Sci.* **7**, 83-102.

Moulin, H. (1986). *Game Theory for the Social Sciences*, 2nd ed., New York: New York Univ. Press.

Rochet, J. C. (1980). "Selection of an Unique Equilibrium Value for Extensive Games with Perfect Information," D.P. Ceremade, Université Paris IX.

Samuelson, L. (1992). "Dominated Strategies and Common Knowledge," *Games Econ. Behav.* **4**, 284-313.

Satterthwaite, M. A. and Sonnenschein, H. (1981). "Strategy-Proof Allocation Mechanisms at Differentiable Points," *Rev. Econ. Studies* **48**, 587-597.