12-2020

Local dominance

Emiliano CATONINI

Jingyi XUE
Singapore Management University, JYXUE@smu.edu.sg

SMU ECONOMICS & STATISTICS

Local Dominance

Emiliano Catonini, Jingyi Xue

December 2020

Paper No. 01-2021

Local dominance

Emiliano Catonini* and Jingyi Xue†

December 22, 2020

Abstract

We define a local notion of weak dominance that speaks to the true choice problems among actions in a game tree and does not necessarily require to plan optimally for the future. A strategy is (globally) weakly dominant if and only if it prescribes a locally weakly dominant action at every decision node it reaches, and in this case local weak dominance is characterized by a (wishful-thinking) condition that requires no forward planning. From this local perspective, we identify form of contingent reasoning that are particularly natural, despite the absence of an obviously dominant strategy (Li, 2017). Following this approach, we construct a dynamic game that implements the Top Trading Cycles allocation under a notion of local obvious dominance that captures a form of independence of irrelevant alternatives.

Keywords: weak dominance, obvious dominance, strategy-proofness, implementation

1 Introduction

Mechanism design has recently been concerned with the simplicity of the game. Experimental and empirical evidence have shown that players often fail to recognize the existence of a weakly dominant strategy, which has traditionally been regarded as the gold standard for the simplicity of the mechanism. As a reaction to this problem, Li (2017) proposed the stronger notion of *obvious dominance*, which players can verify

*Higher School of Economics, ICEF, emiliano.catonini@gmail.com
†Singapore Management University, School of Economics, jyxue@smu.edu.sg.

without any contingent reasoning, provided they have guessed their entire obviously dominant strategy from the start. In this paper, we look for dominance relations among the actions available to a player at a decision node, which players can verify without having a definite plan for the future, provided they are able to do some (form of) contingent reasoning.[1]

Local weak dominance does not rely on the existence and the anticipation of future dominance relations. On the other hand, it still allows for full-blown contingent reasoning. Direct mechanisms have no future to plan for, therefore they cannot be formally ruled out when the chosen notion of simplicity is just the unnecessity of optimal forward planning/folding-back planning. Acknowledging this impossibility, we will not formally discriminate between easy and difficult mechanisms solely based on this idea of simplicity, thus we let local weak dominance span from no or one-period foresight, to perfect foresight in terms of required forward planning. Nonetheless, speaking to the true decision problems faced by players and allowing for a flexible use of tentative continuation plans for verification, local weak dominance sheds new light on why some dynamic mechanisms are easy to play, regardless of the existence of an obviously dominant strategy: some local weak dominance relations are very easy to spot, because they rest on very simple albeit suboptimal continuation plans for the comparison. In Section 2, we perform a dynamic transformation of "guess 2/3 of the average" and we analyze clock auctions as examples of local weak dominance relations that can be easily discovered by coupling the dominant action with just one (even dominated) action for the next period. From a normative perspective, our first recommendation is then to design a game tree where the definition of local weak dominance is everywhere easy to verify.

Next, we exploit our new local perspective to identify forms of contingent reasoning that are particularly natural and simple. First, consider a situation where choosing one action or the other "makes no difference" for the continuation of the game. As an extreme example, suppose that before playing an obviously strategy-proof mechanism, a player can redeem a gift for having signed up. Redeeming the gift cannot be observed by the other players and does not alter the mechanism. Thus, there is no meaningful contingent reasoning to make: the strategies of the opponents

[1] In a related fashion, Pycia and Troyan (2019) refine obvious dominance by distinguishing at the outset between simple histories and not in terms of players' ability to anticipate their choices. We do not refine obvious dominance and do not fix players' forecasting horizon.

are simply irrelevant for the decision problem. Yet, obvious dominance considers players who are always concerned that types and strategies of the opponents may matter,[2] but are then unable to conclude they don't. So, according to obvious dominance, players could refrain from redeeming the gift in fear of somehow altering their prospects for the game. In light of this, our notion of *local obvious dominance* only requires players to distinguish between the contingencies in which the current choice "makes a difference" and the contingencies in which it "makes no difference" for the continuation of the game, under a suitable, very general formalization of when it does not. This is a very rough form of contingent reasoning related to independence of irrelevant alternatives. Then, our player will simply ignore the no-difference category of contingencies, and compare her best and worst outcomes under the other contingencies jointly considered, as in obvious dominance. In Section 4, we construct a dynamic game that implements the TTC allocation in locally obvious dominant actions, whereas in the direct mechanism the weakly dominant strategy is not locally obviously dominant, and no mechanism implements the TTC allocation in obviously dominant strategies (Li 2017, Troyan 2019). A more permissive notion than local obvious dominance would further partition the relevant contingencies according to the next information sets our player may reach, or according to the final outcomes: will I regret this choice given what I will learn tomorrow? or given what I will learn at the best possible terminal node that follows? The formalization of these other notions of dominance is subject for further research.

Local weak dominance works as follows. Action a locally weakly dominates action b if for every continuation plan after action b, there is a continuation plan after a that gives sometimes higher and never lower payoff no matter what the opponents do. We say action a is locally weakly dominant if it locally weakly dominates all other actions. The continuation plan after a that does the job needs not be "optimal" in any sense and can (or *must*) change depending on the continuation plan after b under consideration. This flexibility makes local weak dominance particularly easy to check in many circumstances. In our dynamic version of "guess 2/3 of the average", commiting to a number above 67 at the first round will be clearly worse than moving to the second round and commiting to 67, although this will likely be suboptimal as well; in an English auction, leaving at a price you are still happy to pay is clearly

[2] For this reason, obvious dominance can only be applied in context with (almost) private values.

worse than leaving at the next round, although then one may want to continue; in our dynamic TTC game, the continuation plan that makes pointing to the favourite item dominant is not optimal, rather it imitates the alternative under consideration.

Consider now a game where, if all players choose locally weakly dominant actions whenever they exist, they only reach decision nodes where they do exist. This form of "on-path strategy proofness" is weaker than traditional strategy-proofness, and does not need to be recognized by players in advance, it can be discovered little by little as the game unfolds. In our dynamic "guess 2/3 of the average", if somebody fails to play a locally weakly dominant action, the same complications of the static game arise. In principle, our designer could quit the game after a (detectable) deviation, but this requires otherwise unnecessary commitment power and the unlikely ability to restrict our players' moves.[3]

Suppose now that a player has a locally weakly dominant action "everywhere", in the following weak sense: at every decision node that can be reached if she does play her locally weakly dominant actions. Do her locally weakly dominant actions constitute a weakly dominant (reduced) strategy? The answer is yes. The converse also holds: if a (reduced) strategy is weakly dominant, all the actions it prescribes are locally weakly dominant. This is not surprising, given that we do not impose any upper bound on the foresight players need to spot a local weak dominance relation. Thus, our notion of local weak dominance yields an extensive-form characterization of strategic-form weak dominance. Seen from this angle, our characterization decomposes the problem of finding a weakly dominant strategy into smaller, simpler problems. Concretely, a player does not need to pose herself the problem of finding a weakly dominant strategy: by just recognizing the existence of a locally weakly dominant action at every information set, our player will unknowingly carry out her weakly dominant strategy.

As said, the definition of local weak dominance does not impose any constraint on the amount of forward (albeit tentative) planning that is required for verification. When a player has a locally weakly dominant action everywhere, in a sense, there is no need for any forward planning, because then local weak dominance is characterized by the following (otherwise weaker) condition: given each possible profile

[3]When the designer has the ability to quit the game after a detectable deviation and assign suitable payoffs, the "pruning principle" stated by Li (2017) for obvious strategy-proofness applies also here.

of opponents' strategies, the best possible outcome after choosing action a is never worse and sometimes better than the best possible outcome after choosing action b. This condition requires no forward planning, in that our player does not need to find continuation plans that perform well under all opponents' strategies, but just look for the best outcomes *after* fixing each profile of opponents' strategies. When actions a and b satisfy this condition, we say that *a wishfully dominates b*: it is as if one could adapt the choices to the unknown opponents' strategies. A wishfully dominant action needs not be locally weakly dominant if there is no locally weakly dominant actions at some future information set. Nonetheless, wishful dominance could be a heuristic players use to make a choice and move on (see our TTC game). If this is the case, decomposing the direct revelation mechanism into a sequence of locally weakly dominant choices never comes at the cost of requiring any forward planning, so the recommendation for the designer could be to "decompose as much as possible".

A strategy obviously dominates another strategy if, conditional on reaching an information set where they depart, the former does no worse than the latter even when they are evaluated, *respectively*, under the least favourable combination and the most favourable combination of opponents' strategies and types. In this way, obvious dominance does not require players to recognize that their actions cannot have any influence on the opponents' simultaneous actions, or even on their types. Local weak dominance does. Relatedly, while perfect information games suffice for obvious dominance implementation (Ashlagi and Gontszarowski, 2018; Pycia and Troyan, 2018) this is not true for us:[4] limiting the observability of past actions can be crucial to preserve dominance, and can actually result in simpler decision problems — our dynamic implementation of the TTC allocation is a case in point.[5] On the other hand, obvious dominance does require to have a global plan in mind. We will

[4]We also note that in a game with perfect information and private values, obvious dominance and weak dominance generically coincide: When strategy s_i weakly dominates strategy s_i', given any information set h where s_i and s_i' depart, if there is perfect information and payoffs do not depend on the opponents' types, one can construct a profile of opponents' strategies s_{-i} compatible with h such that (s_i, s_{-i}) and (s_i', s_{-i}) give to player i the maximum and the minimum payoffs she can get with s_i and s_i' after h, and by weak dominance and genericity the first is higher than the second. A formal proof is available upon request.

[5]When strategy s_i weakly dominates strategy s_i', given any information set h where s_i and s_i' depart, if there is perfect information and payoffs do not depend on the opponents' types, one can construct a profile of opponents' strategies s_{-i} compatible with h such that (s_i, s_{-i}) and (s_i', s_{-i}) give to player i the maximum and the minimum payoffs she can get with s_i and s_i' after h, and by weak dominance and genericity the first is higher than the second. A formal proof is available upon request.

show in the English auction that a player who does not engage in contingent reasoning needs to figure out in advance the obviously dominant strategy to move on. So, once we fix a way of reasoning there is a trade-off between the two ideas of simplicity, no contingent reasoning and no global planning.

Local obvious dominance improves this trade-off by only requiring a minimal form of contingent reasoning. Given an ordered action pair (a, b) of our player, given a non-negative constant ε, we say that a set of strategies of the opponents $\overline{S}_{-i}$ is ε-*irrelevant* for (a, b) if for every continuation plan after b, there is a continuation plan after a that gives *exactly* ε-higher payoff no matter what the opponents do within $\overline{S}_{-i}$. This very general definition includes the cases where the choice between a and b makes absolutely no difference, the cases where the only difference is a flow payoff given by the choice itself (as in the gift example), and a variety of similar cases where a and b lead to essentially identical subtrees in all non-trivial scenarios — our TTC game will be a case in point. Then, a locally obviously dominates b if there exists a bipartition of the strategies of the opponents compatible with the information set where one set is ε-irrelevant, and over the other set the following is true: there is a continuation plan after a that always yields a non-lower payoff than the best payoff after b. Although seemingly very weak, the condition that at the end of the game the difference between a and b is always the same amount is very powerful in singling out the scenarios where players do not truly need to plan ahead and do contingent reasoning to compare a and b. In a direct mechanism, given two reports a and b, the set of reports of the opponents where the difference between a and b is constant is typically way too small to leave out a set over which obvious dominance holds. In the direct mechanism for TTC, there is no local obvious dominance between two reported rankings unless they are very similar.[6] In our dynamic mechanism, instead, players will only have to point to their favourite item among the still available ones, and this locally obviously dominates pointing to any other items, because the choice is relevant only when it terminates the game for our player, who then leaves with the

[6]If two rankings r and r' differ only by a swap between two adjacent items a and b, a is preferred to b, and r puts a first, then r obviously dominates r': the set of all reports of the opponents except those that yield a under r and b under r' is 0-irrelevant. It is interesting to note that starting from any reported ranking one can obtain the true ranking with a sequence of such swaps, each corresponding to a locally obvious dominant relation. Nonetheless, the true ranking typically does not dominate the first. Thus, local weak dominance is not transitive. This is actually a desirable property for a relation that captures a limitation of cognitive ability: small steps of comparison are feasible for our player, but one larger step is not.

item she pointed to.

The paper is organized as follows. In Section 2 we revisit "guess 2/3 of the average" and ascending auctions. In Section 3 we construct our notions of dominance and provide the equivalence results. In Section 4, we propose and analyze our dynamic TTC game.

2 Two examples: herding game and ascending auctions

Herding game A shepherd dog has to recall the sheep from the top of the hill for the night. His goal is to maximize the number of sheep that make it all the way down to the sheepfold before falling asleep. Then, by contract, the dog has to guard the sheep from 2/3 of their average altitude. The sheep, instead, want to sleep as close as possible to the dog, but at dusk it gets foggy and they can cannot see where the others are going. Moreover, they are too tired to walk uphill. So they typically stop somewhere on their way down and sleep scattered on the slope. To solve this problem, the dog comes up with the following idea. He first positions himself at altitude 67 (the top of the hill is at altitude 100) and barks. The sheep start moving down towards the dog. Those who stop along the way can't help falling asleep. Those who reach the dog get to see each other. Then the dog moves to altitude 45 and barks again. The game continues in this fashion until the dog reaches altitude 1. Then, if some sheep did not make it to the sheepfold, the dog moves up to his prescribed guarding position.

Take now the viewpoint of a sheep at the top of the hill. It is easy to realize that reaching the dog is a good idea: his guarding position will not be above 67 even in case nobody else moves downhill. Then, the sheep can observe how many others reached the dog. If all sheep reached altitude 67, then it is again easy to decide to reach the dog at altitude 45, because his guarding position will not be higher than that. And so on.

This argument can be formalized with local weak dominance. Consider the path where all sheep arrive to the sheepfold. At any information set along this path, a sheep compares stopping somewhere between her current position and the dog's position with reaching the dog and stopping there for the night. The second strategy does strictly better than the first, no matter what the other sheep do: given that all sheep have already reached our sheep's current altitude, the dog won't guard them from a higher altitude than his current barking position.[7] Therefore, if all sheep use local weak dominance, they will reach the sheepfold.

Furthermore, once all sheep have reach the dog, reaching the dog at the next position is also locally obviously dominant. The reason is that the continuation plans of the other sheep are all completely irrelevant for the choice: once concluded that the dog will not sleep above a given altitude x, moving from altitude $y > x$ to x simply reduces the distance from the dog's final position by $y - x$, therefore all the strategies of the other sheep are $(y - x)$-irrelevant.

Note that the sheep do not have a weakly dominant strategy. If at some information set a sheep realizes that somebody else stopped uphill, reaching the dog's altitude might not be optimal, because the dog might have to guard the sheep from a higher altitude. These information sets, though, could be eliminated from the game: if the dog realizes that not all sheep have reached him, he could wait for the sheep to fall asleep and move directly to his guarding position instead of barking again from a lower altitude. Quitting the game in this way, however, requires the ability of the dog to commit to a suboptimal behavior given his objective function.

Ascending auctions Li (2017) considers an English auction with private values and two bidders who bid one at a time. He shows that leaving the auction when the price surpasses the own valuation is obviously dominant, whereas revealing the own valuation in the second-price sealed-bid auction is only weakly dominant, and this could explain the higher rate of truthful bidding observed experimentally in the English auction with respect to the second price auction (Kagel et al. 1987). We consider instead a clock auction for one indivisible object. At every round, the two bidders decide simultaneously whether to stay or leave. If one leaves and the other stays, the one who stays wins the object at the current price. If they both stay, the price is increased by one unit and the auction goes on. If they both leave, the item is sold at the current price to one of the two bidders selected at random.

[7] The argument works (with weak incentives under some strategies of the other sheep) also if a sheep gets positive utility only when all other sheep are farther from the dog, as in the classic "guess 2/3 of the average" game.

In the clock auction, it is not obviously dominant to leave when the current price is equal to the own valuation. Take the viewpoint of bidder 1 with valuation v_1 when the current price is $p < v_1$. Bidder 1 compares the strategy of leaving when the price reaches v_1 with the strategy of leaving immediately. The lowest payoff our player can get with the first strategy is 0. This occurs when bidder 2 stays in the auction until v_1 included. The highest payoff our player can get by leaving immediately is $v_1 - p$, with probability 1/2. This occurs when bidder 2 leaves immediately too. Since the latter payoff is higher than the former, there is no obvious dominance relation between the two strategies.

However, we argue, it is not hard for bidder 1 to understand that leaving is not a good idea when $p < v_1$, and is instead optimal when $p = v_1$. In our view, when $p < v_1$, what makes the choice simple is not the comparison between leaving and the globally optimal strategy, which a player may not have anticipated yet, but a local comparison between the actions of leaving and staying. Compared to leaving immediately, bidder 1 has clearly nothing to lose from staying once more and leaving at the next round: if bidder 2 is leaving, she will win for sure instead of with probability 1/2 and will pay the same price; else, her payoff will be zero or positive, instead of certainly zero. This argument can be formalized also with local obvious dominance, using precisely this bipartition: if bidder 2 leaves, there is a constant benefit of staying, if bidder 2 stays, the worst that can happen by leaving at the next round is equivalent to the sure payoff of leaving immediately.[8] On the contrary, when $p = v_1$, leaving guarantees a payoff of zero, whereas staying can result in winning the object at price $p = v_1$ or higher, in case the bidder 2 stays now and leaves next. Given these simple considerations, our bidders will leave at the own valuation (if reached) without having to realize that it is a dominant strategy.

In the English auction, the comparison between the obviously dominant strategy of bidding up to v_1 and an alternative strategy of bidding up to $p < v_1$ works as follows. The two strategies depart when the price reaches p and the current winner

[8] Here the relevant feature of local obvious dominance for the argument is simply that it allows for a bipartition of the contingencies. This is a way the local perspective offered by the dynamic game and rough contingent reasoning have a positive interaction: the local problem is so simple that a bipartition yields all the contingent reasoning that is needed. Obvious dominance, not even allowing for a bipartition, forces the player to compare leaving under the contingency where also the other bidder leaves, which yields a positive payoff, with staying under the contingency where the other bidder stays longer, which yields a zero payoff.

is bidder 2. At this point, carrying out the alternative strategy implies that bidder 1 leaves the game with zero payoff. Continuining with the obviously dominant strategy, instead, never gives a negative payoff and sometimes gives a positive payoff. So, the best outcome from the alternative strategy is no better than the worse outcome from the obviously dominant strategy. However, suppose that bidder 1 has not come up yet with the obviously dominant strategy, or has not even done any plan for the future at all. Then, if bidder 1 does not do contingent reasoning, the best she can do is comparing the best outcome after leaving with the worst outcome after staying. The first is zero, while the second can be negative. Therefore, bidder 1 cannot discover her obviously dominant little by little without doing contingent reasoning.

3 Local dominance

We consider a dynamic game with finite horizon and partial and asymmetric observation of an initial move by nature. We model it as a tree, that is, as a set of nodes endowed with a precedence relation $\prec$. At the root, nature chooses an action θ from the set Θ. Then, a finite set I of players sequentially choose actions until a terminal node is reached. Let Z denote the set of terminal nodes. For each player $i \in I$, the set of the nodes where player i moves is partitioned into information sets. Let H_i denote the collection of i's information sets. Each H_i satisfies the standard perfect recall assumptions, therefore it inherits from the game tree the partial order $\prec$. At each information set $h \in H_i$, player i chooses an action from the set A_i^h. A reduced strategy of player i (henceforth, just "strategy") is a map s_i that assigns an action $a_i \in A_i^h$ to each information set $h \in H_i$ that can be reached given the actions assigned to the previous information sets. Let S_i denote the set of strategies of player i, and let $S_{-i} := \Theta \times (\times_{j \neq i} S_j)$ denote the set of strategy profiles of all players excluding i and including nature. For each $s \in \Theta \times (\times_{i \in I} S_i)$, let $\zeta(s)$ denote the induced terminal node. Let $u_i : Z \to \mathbb{R}$ denote the payoff function of player i.

For each information set $h \in H_i$, let $S_i(h)$ and $S_{-i}(h)$ denote the sets of strategies (profiles) that are consistent with h. For each available action $a_i \in A_i^h$, let $S_i(h, a_i)$ denote the set of strategies $s_i \in S_i(h)$ with $s_i(h) = a_i$. For each strategy $s_i \in S_i$, let $H_i(s_i)$ denote the set of information sets of i that are consistent with s_i. Finally, given an information set $h \in H_i$, an action $a_i \in A_i^h$, and a profile $s_{-i} \in S_{-i}(h)$, let

$Z(h, a_i, s_{-i})$ denote the set of terminal histories z such that $z = \zeta(s_i, s_{-i})$ for some $s_i \in S_i(h)$ with $s_i(h) = a_i$.

Obvious dominance (Li, 2017) compares two strategies from the information sets where they depart, i.e., where they prescribe different actions.[9] The comparison is performed between the best future outcome compatible with the dominated strategy and the worst future outcome compatible with the dominating strategy. So, the future outcomes the player focuses on are restricted with the own plans, but not by considering one strategy profile of the opponents at a time. A way to capture the lack of optimal forward planning as a notion of simplicity, instead of Li's lack of contingent reasoning, could restrict the future outcomes with each strategy profile of the opponents but not with an own plan. Accordingly, one could establish that an action dominates another action if, for each strategy profile of the opponents (including nature), the best compatible outcome that follows the dominated action is not better than the worst compatible outcome that follows the dominating action. However, this would be a way too strong notion of dominance, which would not be satisfied even in the English auction. The reason is that taking the worst possible outcomes after an action without any restriction on the own future moves means that the player does not even contemplate protecting herself against those outcomes. Therefore, we let players restrict their view of the possible future outcomes also with their own continuation plans. (We will reconsider the approach of not restricting outcomes with own plans later.) We take a player who has not performed a folding-back planning exercise, and therefore does not possess definite plans on how to capitalize on the current choices. However, our player considers different tentative plans after an action. Suppose that, when comparing two actions, for every tentative plan after the first action, our player always finds a continuation plan after the second action that does always better, no matter what the opponents do. Then, she comes to the conclusion that the second action dominates the first.

We can now introduce our notion of local weak dominance.

Definition 1 *Fix an information set $h \in H_i$. Action $\overline{a}_i \in A_i^h$ locally weakly dominates action $a_i \in A_i^h$ (at h) if for every $s_i \in S_i(h, a_i)$, there exists $\overline{s}_i \in S_i(h, \overline{a}_i)$ such*

[9] Li talks of *earliest* point of departure because he considers "full" strategies; two reduced strategies prescribe actions at the same information set only if they have always prescribed the same actions before.

that

$$\forall s_{-i} \in S_{-i}(h), \quad u_i\left(\zeta(\overline{s}_i, s_{-i})\right) \geq u_i\left(\zeta(s_i, s_{-i})\right), \tag{1}$$

$$\exists s_{-i} \in S_{-i}(h), \quad u_i\left(\zeta(\overline{s}_i, s_{-i})\right) > u_i\left(\zeta(s_i, s_{-i})\right). \tag{2}$$

We say that $\overline{a}_i$ is locally weakly dominant if it locally weakly dominates every other $a_i \in A_i^h$.

When player i compares two actions at h, she must consider how she would continue afterwards. However, to conclude that action $\overline{a}_i$ dominates action a_i, player i does not need to find *one* continuation plan after $\overline{a}_i$ that beats *all* continuation plans after a_i. When there is no weakly dominant strategy, such a continuation plan may not even exist. When the game has a repetitive structure, the convenient comparisons often match each continuation plan after a_i with its imitation after $\overline{a}_i$ — our TTC game will be a case in point. When a_i terminates the game, only one continuation plan after $\overline{a}_i$ is needed, but the easiest to compare may even be dominated itself—leaving at the next stage in the ascending auction is an example.

We say that the game is *on-path strategy-proof* when, for every player, there is a locally weakly dominant action at every information set that is reached if every player chooses her locally weakly dominant actions. For implementation purposes, achieving on-path strategy proofness is as good as strategy proofness and does not require to commit to a suboptimal design off-path — in our herding game the shepherd dog can keeping working towards the goal off-path, although at this point the game has become complicated for the sheep.

The traditional notion of strategy-proofness is based on weak dominance in the strategic form of the game.

Definition 2 *A strategy $\overline{s}_i \in S_i$ weakly dominates strategy $s_i \neq \overline{s}_i$ if*

$$\forall s_{-i} \in S_{-i}, \quad u_i\left(\zeta(\overline{s}_i, s_{-i})\right) \geq u_i\left(\zeta(s_i, s_{-i})\right), \tag{3}$$

$$\exists s_{-i} \in S_{-i}, \quad u_i\left(\zeta(\overline{s}_i, s_{-i})\right) > u_i\left(\zeta(s_i, s_{-i})\right). \tag{4}$$

We say that $\overline{s}_i$ is weakly dominant if it weakly dominates every other strategy.

In static games, local weak dominance coincides with weak dominance. In dynamic games, there is the following relationship: if action a_i is locally weakly dominated at

h, then every strategy s_i that prescribes a_i is weakly dominated[10] by the strategy that prescribes the dominating action $\overline{a}_i$ at h, coincides with the continuation plan that does the job against the continuation plan of s_i afterwards, and coincides with s_i elsewhere. So, if there is a locally weakly dominant action at every information set that follows the choice of locally weakly dominant actions, all the other actions can only be prescribed by weakly dominated strategies, and a weakly dominant strategy emerges: the one that prescribes the locally weakly dominant actions. Conversely, a weakly dominant strategy always prescribes only locally weakly dominant actions. This is because local weak dominance does not necessarily impose to use the flexibility in the choice of the continuation plan, and the one derived from a weakly dominant strategy always does the job. We will formalize this equivalence later.

Now, let us reconsider the approach of not restricting the possible outcomes with own continuation plans. The aim is to construct a notion of dominance that does not even rely on tentative forward planning. As said, taking the worst possible outcome after the candidate dominating action is too restrictive. A natural alternative is to consider, given each strategy profile of the opponents, the *best* possible outcome not just after the candidate dominated action but also after the candidate dominating action.

Definition 3 *Fix an information set $h \in H_i$. Action $\overline{a}_i \in A_i^h$ wishfully dominates action $a_i \in A_i^h$ (at h) if*

$$\forall s_{-i} \in S_{-i}(h), \qquad \max_{z \in Z(h,\overline{a}_i,s_{-i})} u_i(z) \geq \max_{z \in Z(h,a_i,s_{-i})} u_i(z), \tag{5}$$

$$\exists s_{-i} \in S_{-i}(h), \qquad \max_{z \in Z(h,\overline{a}_i,s_{-i})} u_i(z) > \max_{z \in Z(h,a_i,s_{-i})} u_i(z). \tag{6}$$

We say that $\overline{a}_i$ is wishfully dominant if it wishfully dominates every other $a_i \in A_i^h$.

In static games, wishful dominance coincides with weak dominance. In dynamic games, one cannot say that an action $\overline{a}_i$ that wishfully dominates action a_i also locally weakly dominates it. This is because condition (5) provides too much flexibility: given

[10] The converse is not true: a strategy can be weakly dominated even if it does not prescribe any locally weakly dominated action. This is because a player who reasons according to local weak dominance may choose an action that is locally undominated but, for instance, does allow to reach the payoff of an outside option she did not take earlier in the game. This reflects the lack of global planning of our player.

two different $s_{-i}, s'_{-i} \in S_{-i}(h)$, one may need two different continuation plans to reach the two best outcomes after $\overline{a}_i$, which may perform bad against the other strategy of the opponent. It is even possible that every continuation plan after $\overline{a}_i$, under some s_{-i}, induces an outcome that is worse than any outcome that follows a_i. Neglecting that it is typically unfeasible to secure the best possible outcome no matter how the opponents play motivates why we call our notion of dominance "wishful" and not just "optimistic". However, if there are wishfully dominant actions at the information sets that follow $\overline{a}_i$, one can use them to construct a continuation plan that achieves the best outcome under every $s_{-i} \in S_{-i}(h)$. In this case, wishful dominance and local weak dominance coincide.

Now we can state and prove these two characterizations of local weak dominance "everywhere".

Theorem 1 *Fix a strategy $\overline{s}_i$. The following are equivalent:*

D1 $\overline{s}_i$ is weakly dominant;

D2 $\overline{s}_i(h)$ is locally weakly dominant at each $h \in H(\overline{s}_i)$;

D3 $\overline{s}_i(h)$ is wishfully dominant at each $h \in H(\overline{s}_i)$.

Proof We are going to show that D1 implies D2, which implies D3, which in turn implies D1.

Fix a weakly dominant strategy $\overline{s}_i$. We are going to show that is locally weakly dominant at each $h \in H_i(\overline{s}_i)$. Fix $a_i \in A_i^h \setminus \{\overline{s}_i(h)\}$, and $s_i \in S_i(h, a_i)$. Note that $\overline{s}_i \in S_i(h, \overline{s}_i(h))$. Conditions (1) and (2) follow directly from (3) and (4).

Fix an action $\overline{a}_i$ that it is locally weakly dominant at an information set h. We are going to show it satisfies conditions (5) and (6). Fix $a_i \in A_i^h \setminus \{\overline{a}_i\}$ and $s_{-i} \in S_{-i}(h)$. Fix $s_i \in S_i(h, a_i)$ such that

$$\max_{z \in Z(h, a_i, s_{-i})} u_i(z) = u_i(\zeta(s_i, s_{-i})).$$

By local weak dominance, there is $\overline{s}_i \in S_i(h, \overline{a}_i)$ such that

$$u_i(\zeta(s_i, s_{-i})) \leq u_i(\zeta(\overline{s}_i, s_{-i})) \leq \max_{z \in Z(h, \overline{a}_i, s_{-i})} u_i(z),$$

which yields condition (5). Moreover, for some $s'_{-i} \in S_{-i}(h)$,

$$u_i\left(\zeta(s_i, s_{-i})\right) < u_i\left(\zeta(\overline{s}_i, s_{-i})\right) \leq \max_{z \in Z(h, \overline{a}_i, s_{-i})} u_i\left(z\right),$$

which yields condition (6).

Now fix a strategy $\overline{s}_i$ that satisfies conditions (5). and (6) at every $h \in H_i(\overline{s}_i)$. We are going to show that it is weakly dominant.

First we show that for each $h \in H_i(\overline{s}_i)$ and $s_{-i} \in S_{-i}(h)$,

$$\max_{z \in Z(h, \overline{s}_i(h), s_{-i})} u_i\left(z\right) = u_i(\zeta(\overline{s}_i, s_{-i})).$$

If player i does not move anymore after $\overline{s}_i(h)$ at h, $\zeta(s_i, s_{-i})$ is constant across all $s_i \in S_i(h)$ with $s_i(h) = \overline{s}_i(h)$. Hence,

$$u_i(\zeta(\overline{s}_i, s_{-i})) = \max_{z \in Z(h, \overline{s}_i(h), s_{-i})} u_i\left(z\right).$$

Now suppose by induction that for every $h' \succ (h, \overline{s}_i(h))$ and $s_{-i} \in S_{-i}(h')$,

$$u_i(\zeta(\overline{s}_i, s_{-i})) = \max_{z \in Z(h', \overline{s}_i(h'), s_{-i})} u_i\left(z\right).$$

Then, for each $s_{-i} \in S_{-i}(h)$, either i does not move anymore after h and $\overline{s}_i(h)$ given s_{-i}, or let h' be the information set of i that immediately follows. We have

$$\max_{z \in Z(h, \overline{s}_i(h), s_{-i})} u_i\left(z\right) = \max_{a_i \in A_i(h')} \max_{z \in Z(h', a_i, s_{-i})} u_i(z) = \max_{z \in Z(h', \overline{s}_i(h'), s_{-i})} u_i\left(z\right) = u_i(\zeta(\overline{s}_i, s_{-i})),$$

where the second equality uses the induction hypothesis.

Now fix $s_i \neq \overline{s}_i$. Fix $s_{-i} \in S_{-i}$. Let h be the last information set of player i such that $h \in H_i(s_i) \cap H_i\left(\overline{s}_i\right) \cap H_i(s_{-i})$. By condition (5), we have

$$u_i(\zeta(s_i, s_{-i})) \leq \max_{z \in Z(h, s_i(h), s_{-i})} u_i\left(z\right) \leq \max_{z \in Z(h, \overline{s}_i(h), s_{-i})} u_i\left(z\right) = u_i(\zeta(\overline{s}_i, s_{-i})).$$

Moreover, by condition (6), there is $s'_{-i} \in S_{-i}(h)$ such that

$$u_i(\zeta(s_i, s'_{-i})) \leq \max_{z \in Z(h, s_i(h), s'_{-i})} u_i\left(z\right) < \max_{z \in Z(h, \overline{s}_i(h), s'_{-i})} u_i\left(z\right) = u_i(\zeta(\overline{s}_i, s'_{-i})).$$

Therefore, $\overline{s}_i$ weakly dominates s_i. ■

Now we introduce our restriction to the amount of contigent reasoning that is needed to identify the local dominance relation. We first introduce our notion of irrelevance of a set of opponents' strategies.

Definition 4 *Fix an information set $h \in H_i$, an ordered pair of actions $(\overline{a}_i, a_i) \in A_i^h \times A_i^h$ and a constant $\varepsilon \geq 0$. A subset of opponents' strategies $\overline{S}_{-i} \subset S_{-i}(h)$ is ε-irrelevant for $(\overline{a}_i, a_i)$ if for every $s_i \in S_i(h, a_i)$, there exists $\overline{s}_i \in S_i(h, \overline{a}_i)$ such that*

$$\forall s_{-i} \in \overline{S}_{-i}, \quad u_i\left(\zeta(\overline{s}_i, s_{-i})\right) = u_i\left(\zeta(s_i, s_{-i})\right) + \varepsilon.$$

We say that $\overline{S}_{-i}$ is irrelevant for $(\overline{a}_i, a_i)$ if there exists $\varepsilon \geq 0$ such that $\overline{S}_{-i}$ is ε-irrelevant

Our notion of irrelevance refers to an *ordered* pair of actions because the constant difference between the continuation payoffs is required to be positive. Hence, according to our definition, if $\overline{S}_{-i}$ is irrelevant for $(\overline{a}_i, a_i)$, it is not irrelevant for the permutation $(a_i, \overline{a}_i)$ unless $\varepsilon > 0$. What we have in mind is the viewpoint of a player who is wondering whether $\overline{a}_i$ dominates a_i, not the other way round, therefore she can safely focus on the complement of $\overline{S}_{-i}$ if after $\overline{a}_i$ she can get a payoff constantly equal to or bigger than after a_i.

Conditional on the scenario where the opponents are *not* playing a irrelevant strategies, things can be complicated, and in the following notion of dominance we do not demand our player to do further contingent reasoning.

Definition 5 *Fix an information set $h \in H_i$. Action $\overline{a}_i \in A_i^h$ locally obviously dominates action $a_i \in A_i^h$ (at h) if there exists $\overline{S}_{-i}$ with $\emptyset \subseteq \overline{S}_{-i} \subseteq S_{-i}(h)$ such that $\overline{S}_{-i}$ is irrelevant for $(\overline{a}_i, a_i)$, and there exists $\overline{s}_i \in S_i(h, \overline{a}_i)$ such that*

$$\min_{s_{-i} \in S_{-i}(h) \setminus \overline{S}_{-i}} u_i\left(\zeta(\overline{s}_i, s_{-i})\right) \geq \max_{s_i \in S_i(h, a_i)} \max_{s_{-i} \in S_{-i}(h) \setminus \overline{S}_{-i}} u_i\left(\zeta(s_i, s_{-i})\right), \tag{7}$$

$$\max_{s_{-i} \in S_{-i}(h) \setminus \overline{S}_{-i}} u_i\left(\zeta(\overline{s}_i, s_{-i})\right) > \max_{s_i \in S_i(h, a_i)} \min_{s_{-i} \in S_{-i}(h) \setminus \overline{S}_{-i}} u_i\left(\zeta(s_i, s_{-i})\right). \tag{8}$$

We say that $\overline{a}_i$ is locally obviously dominant if it locally obviously dominates every other $a_i \in A_i^h$.

Condition (7) says that there exists a continuation plan after the dominating action $\overline{a}_i$ that ensures to player i at least the best payoff she can get after the dominated action a_i. Condition (8) makes sure that, no matter how player i would continue after a_i, the continuation plan after $\overline{a}_i$ can also yield a strictly higher payoff.

When the definition of local obvious dominance is verified with $\overline{S}_{-i} = \emptyset$, every strategy that prescribes the locally dominated action is obviously dominated[11] by the strategy that prescribes the dominating action $\overline{a}_i$ at h, coincides with the continuation plan that does the job against the continuation plan of s_i afterwards, and coincides with s_i elsewhere. Moreover, an obviously dominant strategy prescribes actions that are locally obviously dominant, provided that we slightly strengthen the original definition of Li (2017) for coherence with our condition (8). Given two strategies $\overline{s}_i, s_i \in S_i$, say that an information set $h \in H_i$ is a *point of departure* if there are $\overline{a}_i, a_i \in A_i^h$ $(\overline{a}_i \neq a_i)$ such that $\overline{s}_i \in S_i(h, \overline{a}_i)$ and $s_i \in S_i(h, a_i)$.[12]

Definition 6 *Strategy $\overline{s}_i$ obviously dominates strategy s_i if for every point of departure h,*

$$\min_{s_{-i} \in S_{-i}(h)} u_i \left(\zeta(\overline{s}_i, s_{-i})\right) \geq \max_{s_{-i} \in S_{-i}(h)} u_i \left(\zeta(s_i, s_{-i})\right), \tag{9}$$

$$\max_{s_{-i} \in S_{-i}(h)} u_i \left(\zeta(\overline{s}_i, s_{-i})\right) > \min_{s_{-i} \in S_{-i}(h)} u_i \left(\zeta(s_i, s_{-i})\right). \tag{10}$$

With this, we can claim the following obvious dominance counterpart of Theorem 1.

Proposition 1 *A strategy is obviously dominant if and only if it only prescribes locally obviously dominant actions with empty sets of irrelevant strategies.*

Proof. Only if. Fix the obviously dominant strategy $\overline{s}_i$. Fix $h \in H_i$ with $\overline{s}_i \in S_i(h)$. Fix $a_i \in A_i^h$ with $a_i \neq \overline{s}_i(h)$. Let

$$s_i' = \arg \max_{s_i \in S_i(h, a_i)} \left(\max_{s_{-i} \in S_{-i}(h) \setminus \overline{S}_{-i}} u_i \left(\zeta(s_i, s_{-i})\right) \right).$$

[11] As for local weak dominance, the converse is not true: a strategy can be obviously dominated although it does not prescribe any locally obviously dominated action.

[12] Given our focus on reduced strategies, every point of departure is an *earliest* point of departure in the sense of Li (2017).

Since $\overline{s}_i$ obviously dominates s_i' and h is a point of departure, by (9) and (10) we have

$$\begin{aligned} \min_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(\overline{s}_i, s_{-i})\right) &\geq \max_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(s_i', s_{-i})\right), \\ \max_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(\overline{s}_i, s_{-i})\right) &> \min_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(s_i', s_{-i})\right), \end{aligned}$$

which yields (7) and (8) with $\overline{S}_{-i} = \emptyset$ by construction of s_i'.

If. Fix a strategy $\overline{s}_i$ such that, for every $h \in H_i$ with $\overline{s}_i \in S_i(h)$, $\overline{s}_i(h)$ is locally obviously dominant with $\overline{S}_{-i} = \emptyset$. Fix strategy $s_i \neq \overline{s}_i$ and a point of departure h. Since $\overline{s}_i(h')$ is locally obviously dominant at every $h' \succeq h$, for every $\overline{s}_i' \in S_i(h, \overline{s}_i(h))$ we have

$$\min_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(\overline{s}_i, s_{-i})\right) \geq \min_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(\overline{s}_i', s_{-i})\right).$$

By local obvious dominance at h, there is $\overline{s}_i' \in S_i(h, \overline{s}_i(h))$ such that, by (7) and (8),

$$\begin{aligned} \min_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(\overline{s}_i', s_{-i})\right) &\geq \max_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(s_i, s_{-i})\right), \\ \max_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(\overline{s}_i', s_{-i})\right) &> \min_{s_{-i}\in S_{-i}(h)} u_i\left(\zeta(s_i', s_{-i})\right), \end{aligned}$$

Altogether, we get (9) and (10). ■

4 Application: Top Trading Cycles

The top trading cycles (TTC) algorithm offers an efficient solution to the problem of trading N items among N agents with ownership of one of the items and a strict preference relation over all items. The algorithm is strategy-proof, in the sense that it is weakly dominant for players to submit their true ranking and let the algorithm do the assignment. However, it is not easy for a player to recognize this. For instance, a player might be tempted to rank an item b above an item a she prefers but believes having slim chances to obtain, because she fears she might miss her chance to get item b while the algorithm insists trying to assign her item a unsuccessfully. To avoid these problems, we design a dynamic game, as opposed to an algorithm, with three simplicity features. First, at each stage players are only asked to name their favourite item among the still available ones, and cannot be assigned anything else than that.[13]

[13] This also has the advantage that players do not have to figure out their entire ranking (at once), but just recognize their top item from a subset at a time. We do not formalize this dimension of

Second, this choice remains secret to the other players, thus our player needs not worry that it may (negatively) affect their future choices. Third, players can rest reassured that whenever an opportunity for trade pops up, it remains intact through time and can be exploited later.

For the implementation of the TTC allocation (and other allocation rules), Bo and Hakimov (2020) propose a dynamic game where players, at each stage, can pick an object from an individualized menu. The TTC allocation emerges as an equilibrium of the game, but the mechanism is not strategy-proof, and it is also not a "simple mechanism" in the sense of Borgers and Li (2019), in that it requires more than first-order belief in the rationality of the opponents. Mackenzie and Zhou (2020) propose similar *menu mechanisms* for a more general class of problems, and show that *private* menu mechanisms achieve strategy-proofness.[14]

Our game works as follows. For each player, there is a "you name it, you get it" repository with the items she can immediately get. At the beginning of the game, in the repository of each player there is only the own item. Then, players name an item. Each player's repository is filled with the items of the players who, directly or indirectly, pointed to our player's item — indirectly means that they named the item of a player who named our player's item, and so on. The players who named an item in their repository leave the game with that item. The players who named an item that was assigned to someone else are asked to name a new available item. The other players wait. At every stage, all players observe the set of still available items, but do not observe the content of their repositories.[15] Note that, as long as a player is in the game, her repository can only increase:[16] all the players who are directly or

simplicity because we consider it orthogonal to the one under investigation in this paper. Bonkoungou and Nesterov (2020) investigate the incentive for truthtelling of players when they only have partial information of their preferences.

[14]Mackenzie and Zhou's mechanism is not designed specifically for TTC, therefore it is naturally more complicated than ours. While our players will only have to name their favourite available outcome, in the menu mechanism players have to pick from menus of allocations, and their final assignment might not be the last one they picked. It is interesting to note, however, that the private menu mechanism achieves strategy proofness in a very similar way to our game, in that the evolution of the menus of allocations proposed to a player only conveys information about the players who left the game, which is precisely all our players can observe.

[15]The unobservability of the repositories may raise transparency concerns. Note however that the classical TTC game requires even more faith in the algorithm: players to reveal immediately *all* their private information.

[16]It is probably a good idea to communicate this explicitly to players, so a player does not need

indirectly pointing to her can only wait for her to complete a cycle.

We first show that naming the favourite available item is a locally weakly dominant action at each $\overline{h} \in H_i$. We are going to use the characterization of local weak dominance with wishful dominance. Let a be player i's favourite item and let b be any other available item. Given each $s_{-i} \in S_{-i}(\overline{h})$, as long as player i is in the game, her choices have no influence on the set of available items, thus on what the other players observe, the information sets they reach, hence the items they name. Suppose first that both a and b are already in player i's repository. Then, player i is strictly better off by naming a and condition (6) is satisfied. Suppose now that s_{-i} is such that item a is or will enter player i's repository (unless player i names b and leaves the game earlier with it or with another item she names next). In this case, by naming a, player i will surely get a and condition (5) is satisfied. Suppose finally that s_{-i} is such that a will be assigned to someone else. Then, whatever item player i can get after naming b, she can also get it after naming a — the repository just keeps growing. Therefore, condition (5) is satisfied.

Now we show that naming the favourite available item is also locally obviously dominant. The 0-irrelevant set is the following: let $\overline{S}_{-i}$ be the set of all $s_{-i} \in S_{-i}(\overline{h})$ such that player i will *not* get a after naming a. In this case, we formalize the idea that player i's prospects for the game are exactly the same as if she had named b. Consider any strategy $\widetilde{s}_i \in S_i(\overline{h}, b)$. Take any $s_i \in S_i(\overline{h}, a)$ such that, for each $h \in H(s_i)$ with $\overline{h} \prec h$, $s_i(h) = \widetilde{s}_i(\widetilde{h})$, where $\widetilde{h}$ is the latest information set compatible with $\widetilde{s}_i$ where the history of available items coincides with, or is a prefix of the one at h, In a nutshell, s_i is the strategy that deviates from $\widetilde{s}_i$ by giving a try at item a at $\overline{h}$, and then "catches up" with $\widetilde{s}_i$ in case the attempt is unsuccesful.

To see that s_i is well-defined, observe the following. First, $\widetilde{h}$ is uniquely defined because all the information sets of i compatible with $\widetilde{s}_i$ where the history of available items is a prefix of the one at h are strictly ordered. There remains to show that $\widetilde{s}_i(\widetilde{h})$ is available at h. For any $s_{-i} \in S_{-i}(h)$, as long as player i is in the game, the history of available items along the path $\widetilde{z}$ induced by $(\widetilde{s}_i, s_{-i})$ is the same as along the path z induced by (s_i, s_{-i}). So $\widetilde{z}$ goes through $\widetilde{h}$, and if player i is still in the game at the round of h, she is still pointing at $\widetilde{s}_i(\widetilde{h})$, which is thus available also at h; if player i

to realize that every opponent (along a sequence) that points to her own item cannot move as long as she is in the game.

leaves at an earlier round, $\widetilde{s}_i(\widetilde{h})$ enters the repository of player i at that round along z as well, so it is available later.

We now show that s_i and $\widetilde{s}_i$ satisfy ??. Fix any $s_{-i} \in \overline{S}_{-i}$. We show that with s_i player i gets the same item she gets with $\widetilde{s}_i$. Let z and $\widetilde{z}$ be the paths induced by (s_i, s_{-i}) and $(\widetilde{s}_i, s_{-i})$. Let h and $\widehat{h}$ be the last information sets of player i along z and $\widetilde{z}$, so with s_i player i gets $s_i(h)$ and with $\widetilde{s}_i$ player i gets $\widetilde{s}_i(\widehat{h})$. Since player i will not get a, we have $h \succ \overline{h}$. We show that $\widehat{h}$ coincides with $\widetilde{h}$ defined as above, so that $s_i(h) = \widetilde{s}_i(\widehat{h})$. As observed, $\widetilde{z}$ goes through $\widetilde{h}$. Then, as long as player i is in the game, the history of available items evolves as along z, so $\widetilde{s}_i(\widetilde{h})$ remains available and i cannot switch to another item. Hence, $\widetilde{h}$ is the last information set along $\widetilde{z}$, thus coincides with $\widehat{h}$.

References

[1] Ashlagi, I. and Y. A. Gonczarowski (2018): "Stable Matching Mechanisms are not Obviously Strategy-Proof", *Journal of Economic Theory*, 177, 405-425

[2] Bo, I and R. Hakimov (2020): "Pick-an-object Mechanisms", working paper.

[3] Bonkoungou, S. and A. Nesterov (2020): "Students incentives under partial information", working paper.

[4] Borgers, T. and J. Li (2019): "Strategically simple mechanisms", *Econometrica*, 87(**6**), 2003-2035

[5] Kagel, J., R. Harstad and D. Levin (1987): "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study", *Econometrica*, 55(**6**), 1275-1304

[6] Li, S. (2017): "Obviously Strategy-Proof Mechanisms", *American Economic Review*, 107**(11)**, 307-352, 3257-87.

[7] Mackenzie, A. and Y. Zhou (2020): "Menu mechanisms", working paper.

[8] Pycia, M. and P. Troyan (2018): "Obvious Dominance and Random Priority", working paper.

[9] Pycia, M. and P. Troyan (2019): "A Theory of Simplicity in Games and Mechanism Design", working paper.

[10] Troyan, P. (2019): "Obviously Strategy-Proof Implementation of Top Trading Cycles", *International Economic Review*, 60**(3)**, 1249-1261.