Mechanism Design with Level-k Types: Theory and an Application to Bilateral Trade

Terri Kneeland*

First Draft: April 17 2013
This Draft: October 24, 2018

Abstract

We develop necessary and sufficient conditions for level-k implementation. These conditions establish a set of level-k incentive constraints that are analogous to Bayesian incentive constraints. We also establish necessary and sufficient conditions for robust level-k implementation (implementation that is robust to different specifications of beliefs about levels of reasoning and to any specification of beliefs about payoffs) that are analogous to ex post incentive constraints. We show that in some environments, level-k (robust level-k) implementation is equivalent to Bayesian (ex post) implementation. We then show, via a bilateral trade application, that this is not a general implication.

*Department of Economics, University College London. Address: Gower Street, London, UK, WC1E 6BT. E-mail: t.kneeland@ucl.ac.uk. I am grateful to Mike Peters for many invaluable discussions and guidance. For very helpful comments I would like to thank Vince Crawford, David Freeman, Yoram Halevy, Li Hao, Eliaz Kfir, Wei Li, Rosemarie Nagel, Ariel Rubinstein, Sergei Severinov, Rani Spiegler, and Alexander Wolitzky. The title of this paper is a play on Wolitzky's (2016) title "Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade".

1 Introduction

Laboratory experiments frequently find that behavior deviates from Nash and Bayesian equilibrium predictions when agents interact in novel environments. Non-equilibrium approaches, like level-k and cognitive hierarchy models, that relax the belief consistency assumptions of equilibrium models have been increasingly used to explain this behavior.[1] This empirical evidence prompts the need for extending the analysis of economic phenomena beyond an equilibrium analysis to other behaviorally plausible solution concepts.

This paper contributes to that end by analyzing mechanism design under the level-k solution concept. In the level-k model, agents anchor their beliefs in a naive model of others' behavior and adjust their beliefs by a finite number of iterated best responses.[2] The model is anchored in the behavior of level 0 types, which is exogenously given and assumed to be uniformly random. Level 1 types engage in one level of reasoning and best respond to level 0 behavior. Level 2 types engage in two levels of reasoning and best respond to level 1 types. And so on, with level k types playing a best response to level $(k-1)$ types. This yields a tractable model of strategic behavior in which agents determine their optimal actions in a finite number of steps. The level-k solution concept relaxes the belief consistency assumption of equilibrium by allowing agents to hold (possibly) inaccurate beliefs about the levels of reasoning of their opponents. Our notion of level-k implementability is identical to the notion of Bayesian implementability, except our solution concept is the level-k solution concept: a social choice rule is level-k implementable if for every profile of payoff types and levels, the actions played under the level-k model lead to outcomes that are consistent with the social choice rule.

[1] For, pioneering works in the literature see Stahl & Wilson (1994; 1995), Nagel (1995), Costa-Gomes et al. (2001), and Camerer et al. (2004). For a recent survey of this literature, see Costa-Gomes et al. (2013).

[2] Formally, the model in this paper most closely resembles Strzalecki (2014) which uses a type space approach to model the level-k solution concept under complete information. This paper adapts that approach to allow for incomplete information. However, this approach closely follows the spirit of the models of Crawford & Iriberri (2007), Crawford et al. (2009), and Crawford (2016) which adapted the level-k models to incomplete information environments.

We first establish general necessary and sufficient conditions for level-k implementation (Propositions 1 and 2). The conditions specify a set of level-k incentive constraints that are analogous to standard Bayesian incentive constraints. The incentive constraints require there to exist a function, for each agent, that maps payoff types to outcomes in such a way that truthfully reporting payoff types is optimal for that agent given everyone else is truthfully reporting their payoff type. Level-k incentive constraints allow these functions to differ across agents while Bayesian incentive constraints require these functions to be the same for all agents. The level-k incentive constraints are thus a weak relaxation of Bayesian incentive constraints. If Bayesian incentive constraints hold, then the level-k incentive constraints also hold. However, it may be possible to ensure that the level-k incentive constraints hold without the Bayesian incentive constraints holding.

The level-k solution concept relaxes the belief consistency assumption of Bayesian equilibrium, which allows the possibility to incentive different agents using different functions. In order to do this, we need to allow agents to send reports that transmit information beyond that of just their payoff types. Rather than focusing on mechanisms where agents truthfully report their payoff types, we consider augmented mechanisms where agents can report their payoff types and their level of reasoning.[3,4] Two agents with the same payoff type and different levels of reasoning may have different beliefs. Hence, augmented mechanisms can be used to incentivize agent's differently.

[3]This mirrors the analysis in Bergemann & Morris (2005). They study implementation in incomplete information environments without the common prior assumption. In this environment, types with the same payoff type may hold different beliefs about the payoff types of others. Thus, a type represents not just the payoff type, but also higher-order beliefs about the payoff types of others. The authors use augmented mechanisms where agents report their types and not just their payoff types.

[4]The revelation principle does not hold in the level-k environment. This is true even if we allow augmented mechanisms where all types, with levels greater than zero, truthfully report their type. This results from 'menu effects', discussed in Crawford et al. (2009). Level 1 agents best respond to Level 0 agents which play uniformly randomly. Thus, level 1 agents are not playing a best response to others truthfully reporting their type, and in fact, depending on the mechanism, may place positive weight on strategies that are never played by agents with levels greater than zero. As such, there may be a role for these additional strategies, those played only by level 0 agents, in satisfying the incentives of level 1 agents.

We then ask how robust are level-k mechanisms to alternative specifications of beliefs? There are two strong belief assumptions embedded in the definition of level-k implementation: (i) beliefs about payoffs are determined by a specific common prior;[5] (ii) beliefs about the levels of reasoning of others are determined as in the level-k model. The level-k model imposes very specific beliefs: a level k type believes her opponent is a level $(k-1)$ type. However, the spirit of limited depths of reasoning is maintained whenever a level k type has any beliefs over lower types: $0, \ldots, (k-1)$. In general, we might allow an agent with level k to hold beliefs over all lower levels. We develop the concept of robust level-k implementation which allows for any beliefs over payoff types and (lower) levels of others. We establish general necessary and sufficient conditions for robust level-k implementation (Propositions 3 and 4). The conditions specify a set of robust level-k incentive constraints that are analogous to standard ex post incentive constraints. Similarly to level-k and Bayesian implementation, there is a gap between robust level-k and ex post incentive constraints. If the ex post incentive constraints hold, then the robust level-k incentive constraints hold. But, it may be possible to ensure that the robust level-k incentive constraints hold without the ex post incentive constraints holding.

The ability to do so depends on the environment. We establish two restrictions on the environment where level-k incentive constraints collapse to Bayesian incentive constraints and robust level-k incentive constraints collapse to ex post incentive constraints. The first is when the social planner is implementing a social choice function (a single-valued rule). And, the second is when the message space is restricted to that of the set of payoff types (a direct mechanism). In both these cases, level-k implementation is equivalent

[5]Mechanisms that are robust to relaxing these strong common knowledge assumptions, typically known as the Wilson doctrine, can insure that a social choice rule will be implemented even if the planner does not know agents' beliefs about the payoffs of others. Much of this literature is due to Bergemann & Morris (2005), who investigate aspects of robust mechanism design (relaxing common knowledge of payoff assumptions) while maintaining the assumption of common knowledge of rationality. We investigate robust implementation that relaxes common knowledge of payoffs under the empirically plausible assumption level-k reasoning.

to Bayesian implementation and robust level-k implementation is equivalent to ex post implementation (Corollary 1 and Proposition 5).

However, in other environments it may be that the level-k and robust level-k incentive constraints hold without the Bayesian or ex post incentive constraints holding. We show that in a bilateral trade environment, ex post efficient trade is both level-k and robust level-k implementable (Proposition 6). This is in obvious contrast to both ex post and Bayesian implementation where there is a conflict between ex post efficiency and incentive compatibility.

There is a growing literature that focuses on behavioral mechanism design.[6] This paper adds to this literature by studying implementation under the level-k model. Four other papers study level-k implementation. To the best of my knowledge, no other paper studies the robustness of level-k implementation to common knowledge of payoff and level-k assumptions. Crawford et al. (2009) looks at setting optimal reserve prices in first and second price auctions when agents are level-k types. Gorelkina (2015) provides a level-k analysis of the expected externality mechanism.

Crawford (2016) revisits Myerson & Satterthwaite's (1983) bilateral trade results under level-k implementation when the message set is restricted to the set of payoff types. Crawford considers two cases: (i) one where levels are unobservable and as such the social planner needs to screen both levels and payoff types (same environment as in this paper); and (ii) one where levels are observable, thus the social planner need only screen payoff types. In the first case, Crawford establishes a parallel result to Proposition 5 in this paper that shows that level-k and Bayesian implementation are equivalent when the message set is restricted to the set of payoff types. And, hence shows the Myerson and Satterthwaite impossibility result for ex post efficient trade holds

[6] Eliaz (2002) studies mechanism design when there is a proportion of 'faulty' agents that fail to act optimally. Glazer & Rubinstein (2012) allow the content and framing of the mechanism to play a role in behavior. de Clippel (2014) studies mechanism design when agents are not rational. Saran (2011) shows that ex post efficient trade can be achieved under bilateral trade when there is a proportion of truthful traders. Wolitzky (2016) investigates mechanism design and bilateral trade when agents are maxmin expected utility maximizers. Glazer & Rubinstein (1998), Eliaz & Spiegler (2006; 2007; 2008), Severinov & Deneckere (2006) study behavioral mechanism design in individual decision problems.

for level-k implementation when the message set is restricted. Crawford also explores what the 'second-best' level-k mechanisms look like in cases where full ex post efficient trade cannot be achieved. In the latter case, Crawford shows that when levels are observable, a setting not explored in this paper, the relationship between level-k and Bayesian implementation is ambiguous and that the Myerson and Satterthwaite impossibility result can break down. Crawford gives a complete Myerson-Satterthwaite-style characterization of the optimal (restricted) mechanism in this case.

de Clippel et al. (forthcoming) establish a set of necessary and sufficient conditions for level-k implementation in a general setting where the social planner aims to implement a single-valued social choice rule. They use a slightly stronger definition of level-k implementation than the one used here - requiring a version of strict implementation where this paper allows indifferences. This implies, in the single-valued choice rule setting, that level-k implementation implies Bayesian implementation (but not necessarily vice versa as it does in our setting). In contrast, we establish a set of necessary and sufficient conditions for the case of a (possibly) multi-valued social choice rule and show that they collapse to Bayesian incentive constraints when the social choice rule is single-valued (Corollary 1). de Clippel et al. also studies level-k implementation under more general level 0 behavior than uniformly random behavior; while we maintain the level 0 assumption of uniformly random behavior throughout this paper.[7]

The previous two papers place restrictions on the implementation problem: Crawford restricts the message space to be equal to the set of payoff types and de Clippel et al. restrict attention to single-valued social choice rules. A general takeaway from these papers is that (when levels are unobservable) level-k implementation is weakly more restrictive than Bayesian implementation. However, the current paper shows that this conclusion arises from the

[7]However, the necessary conditions for level-k and robust level-k implementation provided in this paper are independent of level 0 behavior. And, while the sufficient conditions for level-k and robust level-k implementation are stated for uniform random level 0 behavior, they could be relaxed for weaker assumptions on level 0 behavior following de Clippel et al. (forthcoming).

restrictions on the implementation problem. If the social planner is interested in multi-valued choice rules and willing to use general message spaces, then level-k implementation can be strictly less restrictive than Bayesian implementation. The bilateral trade application is one such example.

This rest of the paper proceeds as follows. Section 2 sets up the general payoff environment and formalizes level-k implementation. Section 3 establishes necessary and sufficient conditions for level-k implementation. Section 4 introduces the concept of robust level-k implementation and establishes necessary and sufficient conditions. Section 5 looks at two examples of special environments that lead to equivalence between level-k and Bayesian implementation and robust level-k and ex post implementation. Section 6 sets up the bilateral trade environment and shows that ex post efficient trade is robust level-k implementable. Omitted proofs can be found in Appendix A.

2 Setup

2.1 General payoff environment

There is a finite set of agents $I = 1, 2, \dots, n$. Agent i's *payoff type* is $\theta_i \in \Theta_i$, where Θ_i is a finite set. There is a compact set of outcomes Y. Each agent has a continuous utility function $u_i : Y \times \Theta \to \mathbb{R}$. Note that we use the notation $X = X_1 \times \cdots \times X_n$ and $X_{-i} = X_1 \times \cdots \times X_{i-1} \times X_{i+1} \times \cdots \times X_N$ for sets $\{X_i\}_{i \in I}$ throughout this paper.

There is a social planner who is concerned with implementing a social choice rule $F : \Theta \to 2^Y \backslash \emptyset$. The planner would like the outcome to be an element of $F(\theta)$ whenever the true payoff type profile is θ.

2.2 Type spaces

We use the framework of a type space in order to formally define agents' beliefs about the payoff types of others. The standard way to do this is to use a Bayesian type space. The set of payoff types along with a common prior over the set of payoff types constitutes a Bayesian type space.

Definition. A **Bayesian type space** $\mathcal{B}$ is a structure $\mathcal{B} = \langle \Theta; \rho \rangle$, where $\rho \in \triangle(\Theta)$.

Given the common prior ρ, each payoff type forms her beliefs by conditioning on the common prior according to Bayes' rule. The belief of an agent with payoff type θ_i about the payoff types of others is given by $\rho(\theta_{-i}|\theta_i) = \frac{\rho(\theta)}{\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_i, \theta_{-i})}$.

Similarly, we use a type space approach to define the level-k model. The level-k type space generates types that differ both by their payoff type and their level of reasoning.

Definition. A $\mathcal{B}$**-based level-k type space** $\mathcal{L}$ is a structure $\mathcal{L} = \langle \mathcal{B}; \bar{k} \rangle$, where $\mathcal{B}$ is a Bayesian type space $\mathcal{B} = \langle \Theta; \rho \rangle$, $\bar{k} \geq 1 \in \mathbb{N}$

Given a $\mathcal{B}$-based level-k type space $\mathcal{L}$, we can define a set of types for each agent as $T_i = \Theta_i \times \{0, 1, \ldots, \bar{k}\}$. An agent's *type* $t_i = (\theta_i, k_i) \in T_i$ is 2-dimensional, representing both her payoff type, θ_i, and her level, k_i. An agent's level represents her level of reasoning - an agent with a level k uses only k steps of reasoning in order to figure out her optimal behavior in any game.[8]

An agent's beliefs about the types of others are determined both by her payoff type and her level. The beliefs of type, $t_i = (\theta_i, k_i)$, about the types of others are determined by the function $\mu_i(t_{-i}|t_i)$:

$$\mu_i(t_{-i}|t_i) = \begin{cases} \rho(\theta_{-i}|\theta_i) & \text{if } k_j = k_i - 1 \; \forall j \neq i \\ 0 & \text{otherwise} \end{cases}.$$

An agent with a level k puts weight only on types that have levels $(k-1)$. This captures the core assumption of the level-k model. An agent's beliefs about the payoff types of other agents are determined by the common prior ρ. Thus, an agent with payoff type θ_i and level k believes that the payoff types of other agents are determined by $\rho(\cdot|\theta_i)$ and that others have level $k-1$.

[8] The bound on the level of reasoning is not necessary, the results in this paper go through if $T_i = \Theta_i \times \{0, 1, 2, \ldots\}$, however bounding the depths of reasoning maintains the finiteness of the type space for simplicity.

We formally call this type space a Bayesian-based level-k type space because beliefs about payoff types are derived from a common prior. We drop this formalism throughout the rest of this paper and refer to these type spaces as simply level-k type spaces.

2.3 Solution concepts

A mechanism specifies an action set for each agent and a mapping between action profiles and outcomes.

Definition. A **mechanism** $\langle M, f\rangle$ consists of a set of actions $M = M_1 \times \cdots \times M_n$ and a function $f : M \to Y$.

Given the payoff environment and (Bayesian or level-k) type space, a mechanism defines a n-agent incomplete information game with action set M_i and payoffs defined by $u_i : Y \times \Theta \to \mathbb{R}$ and $f : M \to Y$ for agent i.

For a given level-k type space, we can define the level-k solution concept. The level-k solution concept imposes that all types, $k \geq 1$, are rational (that is, they play a best response given their beliefs about the actions of other agents) and have consistent beliefs about the actions of other types. Level 0 types are assumed to play uniformly randomly.[9]

Definition. For a given game defined by a mechanism $\langle M, f\rangle$ and level-k type space $\mathcal{L} = \langle \mathcal{B}; \bar{k}\rangle$, a strategy profile $s = s_1 \times \cdots \times s_n$, with $s_i : \Theta_i \times \{1, \ldots, \bar{k}\} \to \triangle(M_i)$ for all $i \in I$, is a **level-k solution** if and only if:

(i) $\int_{m_{-i}\in M_{-i}} \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f((s_i(\theta_i, 1), m_{-i}), \theta) dm_{-i}$

$\geq \int_{m_{-i}\in M_{-i}} \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f((m', m_{-i}), \theta) dm_{-i} \;\; \forall m' \in M_i,\; \theta_i \in \Theta_i,\; i \in I.$

(ii) $\sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f(s_i(\theta_i, k), s_{-i}(\theta_{-i}, k-1)), \theta)$

[9] The behavior of level 0 types is specified outside of the model. Thus, level 0 types do not play a best response to their beliefs (and may play actions that are not a best response to any belief). In fact, there are no restrictions on a level 0's beliefs in a level-k type space and we do not formally define them.

$$\geq \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f(m', s_{-i}((\theta_{-i}, k-1)), \theta) \;\; \forall m' \in M_i,\; \theta_i \in \Theta_i,$$
$k \in \{2, \ldots, \bar{k}\}$, $i \in I$

The level-k solution can be calculated recursively given the behavior of level 0 types. Level 1's actions are a best response to level 0's actions (condition (i)). Level 2's actions are a best response to level 1's actions, and so on (condition (ii)).

2.4 Implementation

A social choice rule, F, is level-k implementable if there exists a mechanism and a level-k solution that achieves F for every message sent. The formal definition is given below.

Definition. A social rule is **level-k implementable** on $\mathcal{L}$ if there exists a mechanism $\langle M, f\rangle$ and a message profile $m_i : T_i \to \triangle(M_i)$ for all $i \in I$, such that $m = m_1 \times \cdots \times m_n$ is a level-k solution and m achieves F: $f(m(\theta \times \hat{k})) \in F(\theta)$ for all $\theta \times \hat{k} \in \times \left\{\Theta_i \times \{1, \ldots, \bar{k}\}\right\}_{i\in I}$.

First, notice that our notion of level-k implementability does not require the mechanism to satisfy the social choice rule for level 0 types. This is for two reasons. The first is theoretical, level 0 agents are non-strategic and play all actions randomly, hence the social planner cannot incentivize their behavior.[10] The second motivation is empirical. While there is mixed support, the estimated frequency of level 0 types is typically small (e.g. Arad & Rubinstein 2012; Costa-Gomes et al. 2001; Costa-Gomes & Crawford 2006; Brocas et al. 2014). Thus we interpret the existence of level 0 types in the model, but not in our implementability requirement, as types that exist only in the minds of the other types.

Second, notice that our notion of level-k implementation does not require knowledge of the *actual* distribution of types (and hence levels). This is because implementation requires that the outcome be consistent with the social

[10] If this type of behavior is a concern we should consider an alternative form of implementability. See Eliaz (2002) for one such possibility - the social planner tries to minimize the deviations from the social choice rule.

choice rule for *all* type profiles and hence does not depend upon the distribution of types.[11]

We will be interested in comparing level-k implementation with that of Bayesian implementation. We define Bayesian implementation for completeness. In the below definition, we incorporate the results of the revelation principle, and hence define Bayesian implementation under a direct mechanism. Condition (ii) below states the standard Bayesian incentive constraints which will be contrasted with the level-k incentive constraints developed in the next section.

Definition. A social choice rule F is **Bayesian implementable** on $\mathcal{B}$ if there exists a mechanism $\langle \Theta, f \rangle$ such that the following conditions hold:

(i) $f(\theta) \in F(\theta) \; \forall \theta \in \Theta$

(ii) $\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f(\theta), \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f(\theta', \theta_{-i}), \theta) \; \forall \theta' \in \Theta_i$, $\theta \in \Theta$, $i \in I$

3 Necessary and sufficient xonditions for level-k implementation

This section establishes necessary and sufficient conditions for level-k implementation that take the form of incentive constraints. The level-k incentive constraints are analogous to the standard Bayesian incentive constraints. Proposition 1 states the necessary conditions and Proposition 2 states the sufficient conditions for level-k implementation.

Proposition 1. *(Necessary Conditions) Let F be a social choice rule. Let $\mathcal{B}$ be a Bayesian type space and let $\mathcal{L} = \langle \mathcal{B}; \bar{k} \rangle$ be a ($\mathcal{B}$-based) level-k type space with $\bar{k} \geq 2$. If F is level-k implementable, then there exists a function $f^i : \Theta \to Y$ for each $i \in I$, such that the following conditions hold:*

[11] This is not true for all mechanism design objectives. For example, it would not be true if the goal of the planner was to maximize expected revenue. If different levels (and payoff types) play different actions with different revenue consequences, then expected revenue will depend upon the actual distribution of both payoffs and levels.

(i) $f^i(\theta) \in F(\theta)\ \forall \theta \in \Theta,\ i \in I$

(ii) $\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f^i(\theta), \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f^i(\theta', \theta_{-i}), \theta)\ \forall \theta' \in \Theta_i,$ $\theta \in \Theta,\ i \in I$

Proof:

Suppose that the social choice rule F is level-k implementable. Then there exists some mechanism $\langle M, g \rangle$ and a function $m_i : T_i \to \triangle(M_i)$ for each $i \in I$ such that $m = m_1 \times \cdots \times m_N$ is a level-k solution and level-k achieves F.

Consider the behavior of an agent i with type $t_i = (\theta_i, 2)$. Then, it must be true that:

$$\begin{aligned} &\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) \cdot u_i(g(m_i(\theta_i, 2), m_{-i}(\theta_{-i}, 1)), \theta) \\ &\qquad \geq \sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) \cdot u_i(m', m_{-i}(\theta_{-i}, 1)), \theta) \quad \forall m' \in M_i \end{aligned} \tag{1}$$

Note that we use the notation $(\theta, k) = ((\theta_1, k), \ldots, (\theta_n, k))$ throughout this talk.

Define $f^i(\theta) = g(m_i(\theta_i, 2), m_{-i}(\theta_{-i}, 1))$ for all $\theta \in \Theta$ and for all $i \in I$.

Condition (ii) holds from (1). Condition (i) holds by definition of $\langle M, g \rangle$ and m level-k implementing F.

□

Notice that the necessary conditions for level-k implementation are generated only from the incentive requirements of a level 2 agent. For this reason, the necessary conditions are independent of the exact specification of level 0 behavior. But, this is also why the conditions in Proposition 1 are not sufficient conditions. We need to also incentive other levels, and specifically, level 1 agents.

The need to incentivize level 1 agents is why the revelation principle fails under level-k implementation. Level 1 types believe that others are placing some weight on all strategies. Thus adding strategies into the mechanism that are not played by any type with a non zero level can still affect the incentives of level 1 agents.

Regardless of the failure of the revelation principle, it is still possible to provide general sufficient conditions for level-k implementation in a widely applicable special case - that of independent private values. Define an environment of private values to be one where utility functions are such that $u_i : Y \times \Theta_i \to \mathbb{R}$ for all $i \in I$. And, define an environment of independent values to be one where the Bayesian type space, $\mathcal{B} = \langle \Theta; \rho \rangle$, is such that $\rho = \prod_i p_i$ for some $p_1 \times \cdots \times p_n \in \triangle(\Theta_1) \times \cdots \times \triangle(\Theta_n)$. The proof of Proposition 2 builds on the mechanism developed in de Clippel et al. (forthcoming) to agents to report both payoff types and levels.

Proposition 2. *(Sufficient Conditions) Let F be a social choice rule and let the environment be one of independent private values. Let $\mathcal{B}$ be a Bayesian type space and let $\mathcal{L} = \langle \mathcal{B}; \bar{k} \rangle$ be a ($\mathcal{B}$-based) level-k type space. F is level-k implementable if there exists a function $f^i : \Theta \to Y$ for each $i \in I$, such that the following conditions hold:*

(i) $f^i(\theta) \in F(\theta)\ \forall \theta \in \Theta,\ i \in I$

(ii) $\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f^i(\theta), \theta_i) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) u_i(f^j(\theta', \theta_{-i}), \theta_i)\ \forall \theta' \in \Theta_i$, $\theta \in \Theta,\ j \in I,\ i \in I$

PROOF:

Consider the following mechanism where the message space for agent i is equal to $M_i = \Theta_i \times \{0, 1, \ldots \bar{k}\} \times [-1, 1]$ and consists a report of her payoff type $\theta_i \in \Theta_i$, level $k_i \in \{0, 1, \ldots, \bar{k}\}$, and a real number, $z_i \in [-1, 1]$.

Let the indicator function $I_i(z) : Z \to \{0, 1\}$ be defined as follows:

$$I_i(z) = \begin{cases} 1 & \text{if } z_i = m_j z_j \text{ for some } m_j \in \mathbb{Z},\ \forall j \in I \\ 0 & \text{otherwise} \end{cases}$$

Define $\tilde{\theta}_i : M \to \Theta_i$ and $\tilde{k}_i : M \to \{0, 1, \ldots, \bar{k}\}$ in the following way. For a given message profile $m = (\theta, k, z)$, if $I_i(z) = 1$ the planner takes the reports as given and sets $\tilde{\theta}_i(m) = \theta_i$ and $\tilde{k}_i(m) = k_i$; otherwise the planner sets $\tilde{\theta}_i(m)$ to some randomly chosen Θ_i according to the prior p_i and $\tilde{k}_i(m) = 0$. The planner then assigns outcomes based on the reports $\tilde{\theta} \times \tilde{k}$ according to the mechanism $g : \Theta \times \{0, \ldots, \bar{k}\}^n \to Y$ defined by

$$g(\theta \times \hat{k}) = \begin{cases} f^i(\theta) & \text{if } k_j = k_i - 1 \text{ for all } j \neq i \in I \\ f^1(\theta) & \text{otherwise} \end{cases}.$$

Consider an agent i with payoff θ_i and level 1. She believes that all other agent's are level 0 agents that are sending messages uniformly randomly. Thus, the probability that $I_j(z) = 1$ is zero for any $z_i \in [-1, 1]$. Hence, our level 1 agent believes that the planner will almost surely use a payoff type for player j that is picked at random according to the prior p_j. Further, if our level 1 agent sends a non-zero report, $z_i \neq 0$, she will expect the planner to disregard her payoff type report and choose randomly according to p_i with probability 1. However, if our level 1 agents sends a zero report, $z_i = 0$, she will expect the planner to use her payoff type as reported.

Thus, if she sends the message (θ_i, k_i, z_i) with $z_i = 0$ and $k_i = 1$, she will expect to receive the following lottery over outcomes

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i} | \theta_i) \cdot f^i(\theta).$$

If she sends the message (θ_i, k_i, z_i) with $z_i = 0$ and $k_i \neq 1$, she will expect to receive the following lottery over outcomes

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i} | \theta_i) \cdot f^j(\theta)$$

for some $j \in I$.

And, if she sends the message (θ_i, k_i, z_i) with $z_i \neq 0$, she will expect to receive the following lottery over outcomes

$$\sum_{\theta \in \Theta} \rho(\theta) \cdot f^1(\theta).$$

Her expected utility from sending (θ_i', k_i, z_i) with $z_i = 0$ when her payoff type is θ_i is then given by

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) \cdot u_i(f^j(\theta', \theta_{-i}), \theta_i)$$

for some $j \in I$.

By condition (ii) we have that

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) \cdot u_i(f^i(\theta), \theta_i) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) \cdot u_i(f^j(\theta', \theta_{-i}), \theta_i)$$

for all $\theta' \in \Theta_i$ and $j \in I$.

It must then also be true that

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i}|\theta_i) \cdot u_i(f^i(\theta), \theta_i) \geq \sum_{\theta \in \Theta} \rho(\theta) \cdot u_i(f^1(\theta), \theta_i).$$

Thus, for agent i with payoff type θ_i and level 1, reporting $(\theta_i, 1, 0)$ is a best response.

We now prove that an agent with payoff type θ_i and level k will send the message $(\theta_i, k_i, 0)$ by induction on the following statement: Let $k \geq 1$ and assume that if for all $l \in \{1, \ldots, k-1\}$, $\theta_j \in \Theta_j$, and $j \in I$ an agent j with payoff type θ_j and level l will report $(\theta_j, l, 0)$, then an agent i with payoff type θ_i and level k will report $(\theta_i, k, 0)$.

The result is true for $k = 1$ by the above argument. No, consider an agent i with payoff type θ_i and level $k \in \{2, \ldots, \bar{k}\}$. She expects that all other agents will be sending reports $z_j = 0$, $k_j = k - 1$, and truthfully reporting their payoff type. Thus, she expects that the social planner will always take their payoff and level reports as given.

Thus, if she sends the message $(\theta_i, k, 0)$ she will expect to receive the following lottery over outcomes

$$\sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot f^i(\theta).$$

If she sends the message $(\theta_i, k_i, 0)$ with $k_i \neq k$, she will expect to receive the following lottery over outcomes

$$\sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot f^j(\theta)$$

for some $j \in I$.

And, if she sends the message (θ_i, k_i, z_i) with $z_i \neq 0$, she will expect to receive the following lottery over outcomes

$$\sum_{\theta\in\Theta} \rho(\theta)\cdot f^j(\theta)$$

for some $j \in I$.

By the same argument above, condition (ii) then implies that our agent will send the report $(\theta_i, k, 0)$.

Therefore, if we define $m_i(\theta_i, k_i) = (\theta_i, k_i, 0)$ for all $\theta_i \in \Theta_i$, $k_i \in \{1, \ldots, \bar{k}\}$ and $i \in I$, then m is a level-k solution and achieves F by condition (i). Hence, F is level-k implementable.

□

Condition (ii) in Proposition 2 generates a set of level-k incentive constraints that generalize the standard Bayesian incentive constraints. The difference between the two is that the level-k incentive constraints can be satisfied with a different function, f^i, for each agent, whereas the Bayesian incentive constraints must hold using the same function, f, for all agents. The relaxation of the cross-player restriction ($f^1 = \cdots = f^n$) arises because of the relaxation of consistent beliefs under the level-k model. A level 3 agent believes she is facing level 2 agents while a level 2 agent believes she is facing level 1 agents. Thus, all agents' incentive constraints can be satisfied by different f functions: a level 3 agent for player i with payoff type θ_i thinks she will receive $f^i(\theta)$ when playing against level 2 agents with payoff type profile θ_{-i}, while a level 2 agent for player j with payoff type θ_j thinks she will receive $f^j(\theta)$ when playing against level 1 agents with payoff type profile θ_{-j}.

Whether the cross-player restriction imposed under Bayesian implementation has bite depends on the environment. In Section 5 we consider two special environments which lead to the cross-player restrictions being automatically imposed. In these environments, there is an equivalence between level-k and Bayesian implementation.

4 Robustness of level-k mechanisms

4.1 Robust level-k implementation

This section shows that we can design a level-k mechanism that is robust to relaxing the two strong belief assumptions present in the level-k type space. One assumption is that beliefs about payoffs are determined by a specific common prior. The other assumption is that beliefs about levels are determined as in the level-k model. Specifically, if an agent is level k, she believes that others have levels exactly equal to $(k-1)$. In general, we might allow an agent with level k to hold beliefs over all lower levels.[12] As long as a level k type only puts weight on lower levels, the spirit of limited depth of reasoning is maintained with each type being able to calculate her optimal action recursively, in a finite number of steps.

In this section, we generalize our type space and solution concept to the limited depth of reasoning (LDoR) concept to allow for arbitrary beliefs about

[12]Cognitive hierarchy models relax the level-k belief structure in this way. In the cognitive hierarchy model, a level k type has beliefs over all lower levels determined by a conditional Poisson distribution. See Camerer et al. (2004) for specifics.

both payoff types and depths of reasoning of others. We then find necessary (Proposition 3) and sufficient conditions (Proposition 4) for robust level-k implementation.

The following definition of a limited depth of reasoning (LDoR) type space generalizes the level-k type space. The LDoR type space allows an agent to hold any belief over payoff types and lower levels. This approach is based on Strzalecki (2014) who develops the framework for games of complete information. We expand the framework here to allow for incomplete information.

Definition. A **limited depth of reasoning type space (LDoR type space)** is a type space $\mathcal{L}^{LDoR} = \langle (T_i, k_i, \theta_i, b_i)_{i=1,\ldots,n}, \bar{k} \rangle$ such that T_i is a finite set for all $i \in I$, $k_i : T_i \to \{0, \ldots, \bar{k}\}$, $\theta_i : T_i \to \Theta_i$, and $b_i : T_i \to \triangle(T_{-i})$ such that

$$b_i(t_i)(\{t_{-i} \in T_{-i} | k_{-i}(t_{-i}) < k_i(t_i)\}) = 1 \quad \forall t_i \in T_i \text{ with } k_i(t_i) > 0 \text{ and } \forall i \in I.$$

The set T_i is a set of types for each player. As in the level-k type space, a player's type represents both her payoff type, θ_i, and her level, k_i. We abuse notation here and also let θ_i and k_i be functions that map types to payoff types and levels, respectively. Thus, for a type, t_i, her payoff type is given by $\theta_i(t_i)$ and her level is given by $k_i(t_i)$. In addition to payoff type and level, a player's type also represents her beliefs about the types of others. The belief function, b_i, specifies, for each type, what she believes about the types of others. Beliefs have the property that each type only puts positive weight on types that have strictly lower levels. This captures the core assumption of the limited depth of reasoning literature and ensures that agents can calculate their optimal actions in a recursive fashion with a finite number of steps given the behavior of level 0 types. We place no assumptions on an agent's beliefs about payoff types.

Given the definition of an LDoR type space, we can define the analogous solution and implementation concepts: the LDoR solution and LDoR implementation.

Definition. For a given game defined by a mechanism $\langle M, f \rangle$ and type space

$\mathcal{L}^{LDoR} = \langle (T_i, k_i, \theta_i, b_i)_{i=I}, \bar{k} \rangle$, a strategy profile $s = s_1 \times \cdots \times s_n$,with $s_i : T_i \to \triangle(M_i)$ for all $i \in I$, is th**e LDoR solution** if and only if:

(i) $s_i(t_i) \sim U[M_i]$ for all $t_i \in \{T_i | k_i(t_i) = 0\}$

(ii) $\sum\limits_{t_{-i} \in T_{-i}} b_i(t_{-i}|t_i) u_i(f(s(t), \theta(t)) \geq \sum\limits_{t_{-i} \in T_{-i}} b_i(t_{-i}|t_i) u_i(f(m, s_{-i}(t_{-i})), \theta(t))$
$\forall\ m \in M_i$, $t_i \in T_i$ with $k_i(t_i) \geq 1$, $i \in I$.

The LDoR solution is similar to the level-k solution. It specifies that all level 0 types play uniformly randomly. And, all types with levels at least 1 play a best response given their beliefs about the types of others and the actions of those types under s.

Definition. A social choice rule is **LDoR implementable** on $\mathcal{L}^{LDoR} = \langle (T_i, k_i, \theta_i, b_i)_{i=I}, \bar{k} \rangle$ if there exists a mechanism $\langle M, f \rangle$ and a message profile $m_i : T_i \to \triangle(M_i)$ for all $i \in I$, such that $m = m_1 \times \cdots \times m_n$ is an LDoR solution and $f(m(t)) \in F(\theta(t))$ for all $t \in \{t \in T | k_i(t_i) \geq 1, \forall i \in I\}$.

Ultimately, we are interested in robust implementation - whether there exists a single mechanism that can implement the social choice rule under *any* LDoR type space. We call this robust level-k implementation.

Definition. Fix a $\bar{k}$. A social rule F is **robust level-k implementable** if there exists a mechanism $\langle M, f \rangle$ such that for any LDoR type space $\mathcal{L}^{LDoR} = \langle (T_i, k_i, \theta_i, b_i)_{i=1,\ldots,n}, \bar{k} \rangle$, F is LDoR implementable under $\langle M, f \rangle$.

We will be interested in comparing robust level-k implementation with that of ex post implementation. We define ex post implementation for completeness. In the below definition, we incorporate the results of the revelation principle, and hence define ex post implementation for a direct mechanism. Condition (ii) below states the standard ex post incentive constraints which will be contrasted with the robust level-k incentive constraints developed in the next subsection.

Definition. A social choice rule F is **ex post implementable** if there exists a mechanism $\langle \Theta, f \rangle$ such that

(i) $f(\theta) \in F(\theta)\ \forall \theta \in \Theta$

(ii) $u_i(f(\theta), \theta) \geq u_i(f(\theta', \theta_{-i}), \theta)\ \forall \theta' \in \Theta$, $\theta \in \Theta$, $i \in I$

4.2 Necessary and sufficient conditions for robust level-k implementation

Propositions 3 and 4 give the necessary and sufficient conditions for robust level-k implementation. The necessary and sufficient conditions define a set of robust level-k incentive constraints that are analogous to ex post incentive constraints. The relationship between robust level-k implementation and ex post implementation mirrors that of the relationship between level-k and Bayesian implementation. If a social choice rule is ex post implementable, it is robust level-k implementable. But, as we will see in Section 6, there are social choice rules that are robust level-k implementable that are not ex post implementable.

Proposition 3. *(Robust Necessary Conditions) Let F be a social choice rule. Let $\bar{k} \geq 2$. If F is robust level-k implementable, then there exists a function $f^i : \Theta_i \to Y$ for each $i \in I$, such that the following conditions hold:*

- **(i)** $f^i(\theta) \in F(\theta)\ \forall \theta \in \Theta, \forall i \in I$
- **(ii)** $u_i(f^i(\theta), \theta) \geq u_i(f^i(\theta', \theta_{-i}), \theta)\ \forall \theta' \in \Theta_i,\ \theta \in \Theta,\ i \in I$

Proof:

Suppose that the social choice rule F is robust level-k implementable. Then there exists some mechanism $\langle M, g\rangle$ that LDoR implements F for any LDoR type space.

Consider the LDoR type space $\mathcal{L}^* = \langle (T_i, k_i, \theta_i, b_i)_{i\in I}, \bar{k}\rangle$, with $T_i = \{(t_i^{0,\theta})_{\theta\in\Theta}, (t_i^{1,\theta})_{\theta\in\Theta}, (t_i^{2,\theta})_{\theta\in\Theta}\}$, $\theta_i(t_i^{k,\theta}) = \theta_i$, $k_i(t_i^{k,\theta}) = k$, and $b_i(t_i^{l,\theta})(t_{-i}^{l-1,\theta}) = 1$ for all $\theta \in \Theta$, $k \in \{0,1,2\}$, $l \in \{1,2\}$, and $i \in I$. Where we use the notation: $t^{k,\theta} = t_1^{k,\theta} \times \cdots \times t_n^{k,\theta}$ and $t_{-i}^{k,\theta} = t_1^{k,\theta} \times \cdots \times t_{i-1}^{k,\theta} \times t_{i+1}^{k,\theta} \cdots \times t_n^{k,\theta}$

Since F is robust level-k implementable there is a LDoR solution m such that:

(1) $g(m(t_i^{2,\theta}, t_{-i}^{1,\theta})) \in F(\theta)$

(2) $u_i(g(m(t_i^{2,\theta}, t_{-i}^{1,\theta})), \theta) \geq u_i((m_i(g_i^{2,\theta'}), m_{-i}(t_{-i}^{1,\theta})), \theta)\ \forall \theta' \in \Theta_i$

Define $f^i(\theta) = g(m(t_i^{2,\theta}, t_{-i}^{1,\theta}))$ for all $\theta \in \Theta$ and $i \in I$.

Therefore conditions (i) and (ii) hold by conditions (1) and (2) respectively.

□

Notice that similarly to Proposition 1, Proposition 3 does not depend on the level 0 assumption. This is because the necessary conditions are generated from the incentive constraints of level 2 agents. This means that the necessary conditions hold regardless of the assumptions made about level 0 behavior.

The need to incentivize other agents, beyond level 2, is also why the conditions in Proposition 3 are not sufficient conditions. Similar to the case for level-k implementation, we find a set of sufficient conditions for robust level-k implementation for private value environments. The proof follows analogously to that of Proposition 2 and can be found in Appendix A .

Proposition 4. *(Robust Sufficient Conditions) Let F be a social choice rule. Let the environment be one of private values. F is robust level-k implementable if there exists a function $f^i : \Theta_i \to Y$ for each $i \in I$, such that the following conditions hold:*

(i) $f^i(\theta) \in F(\theta)\ \forall \theta \in \Theta, \forall i \in I$

(ii) $u_i(f^i(\theta), \theta_i) \geq u_i(f^j(\theta', \theta_{-i}), \theta_i)\ \forall \theta' \in \Theta_i, \theta \in \Theta,\ j \in I, i \in I$

Condition (ii) in Proposition 4 generates a set of robust level-k incentive constraints that generalize ex post incentive constraints analogously to the relationship between the level-k and Bayesian incentive constraints. The difference between the two is that the robust level-k incentive constraints can be satisfied with a different function, f^i, for each agent, whereas the ex post incentive constraints must hold using the same function, f, for all agents. Whether the cross-player restrictions imposed under ex post implementation have bite depends on the environment. In the next section we consider two special environments which lead to the cross-player restrictions being automatically imposed. In these environments, there will be an equivalence between robust level-k and ex post implementation.

5 Special environments

In this section we look at two special special environments. In the first we restrict attention to social choice functions (single-valued social choice rules) and in the second we restrict attention to mechanisms where the message set is equal to the set of payoff types. These special environments are the environments studied in de Clippel et al. (forthcoming) and Crawford (2016) respectively. In both of these cases, we show that the level-k incentive constraints collapse down Bayesian incentive constraints. This establishes parallel results to those found in de Clippel et al. and Crawford: that level-k and Bayesian implementation are equivalent in these special environments. We further show that robust level-k and ex post implementation are equivalent in these special environments.

Corollary 1 formalizes this result for the restriction to social choice functions.

Corollary 1. *Let F be a social choice function and the environment be one of independent private values. Let $\mathcal{B}$ be a Bayesian type space and let $\mathcal{L} = \langle \mathcal{B}; \bar{k} \rangle$ be a ($\mathcal{B}$-based) level-k type space with $\bar{k} \geq 2$. The following results hold*

(i) *F is level-k implementable if and only if it is Bayesian implementable*

(ii) *F is robust level-k implementable if and only if it is ex post implementable.*

PROOF:

Part (i): (⇒) From Proposition 1, there exists functions $f^i : \Theta \to Y$, for all $i \in I$ such that conditions (i) and (ii) in Proposition 1 hold. Since F is a social choice function, it must be that $f^1(\theta) = \cdots = f^n(\theta) = F(\theta)$ for all $\theta \in \Theta$. Therefore, it follows that F is Bayesian implementable using the mechanism $\langle \Theta, F \rangle$.

(⇐) If F is Bayesian implementable there exists a mechanism $\langle \Theta, f \rangle$ such that the conditions (i) and (ii) in the definition of Bayesian implementability hold. Set $f^i = f$ for all $i \in I$. The conditions of Proposition 2 are satisfied, thus F is level-k implementable.

Part (ii): ($\Rightarrow$) From Proposition 3, there exists functions $f^i : \Theta \to Y$, for all $i \in I$ such that conditions (i) and (ii) in Proposition 3 hold. Since F is a social choice function, it must be that $f^1(\theta) = \cdots = f^n(\theta) = F(\theta)$ for all $\theta \in \Theta$. Therefore, it follows by definition that F is ex post implementable using the mechanism $\langle \Theta, F \rangle$.

($\Leftarrow$) If F is ex post implementable there exists a mechanism $\langle \Theta, f \rangle$ such that conditions (i) and (ii) in the definition of ex post implementability hold. Set $f^i = f$ for all $i \in I$. The conditions of Proposition 4 are satisfied, thus F is robust level-k implementable.

□

Recall that the mechanisms we construct in the proofs of Propositions 2 and 4 are not direct mechanisms (where the set of messages is restricted to be the set of payoff types). Instead, we construct augmented mechanisms where we allow agents with the same payoff type but different levels of reasoning to play different actions. Restrictions to direct mechanisms arise naturally under Bayesian and ex post implementation because the revelation principle holds. However, the revelation principle does not hold under level-k implementation. This suggests that we may not be able to restrict the message space to the set of payoff types without loss of generality. The following proposition demonstrates that restricting the message space to the set of payoff types has the same effect as restricting our social choice rule to a social choice function - the level-k (robust level-k) incentive constraints collapse down to the Bayesian (ex post) incentive constraints. In other words, if we restrict the message space to be the set of payoff types, a social choice rule will be level-k implementable if and only if it is Bayesian implementable. And, it will be robust level-k implementable if and only if it is ex post implementable.

In order to show this we need to assume a richness condition on the environment, Assumption A*:

A*: For every $i \in I$ and for any two payoff types $\theta_i \neq \theta_i' \in \Theta_i$ there exists a $\theta_{-i} \in \Theta_{-i}$ such that $F(\theta_i, \theta_{-i}) \cap F(\theta_i', \theta_{-i}) = \emptyset$.

Assumption A* is likely to hold in many environments. For example, as we'll see in Section 6, it holds in the bilateral trade environment as long as for every two values of the buyer $v < v' \in V$ there exists a cost for the seller, $c \in C$, that falls between, $v \leq c \leq v'$. And, similarly for the seller. Proposition 5 establishes the result.

Proposition 5. *Let F be a social choice rule. Let the environment be one of independent private values where A^* holds. Let $\mathcal{B}$ be a Bayesian type space and let $\mathcal{L} = \langle \mathcal{B}; \bar{k} \rangle$ be a ($\mathcal{B}$-based) level-k type space with $\bar{k} \geq 2$. The following results hold:*

(i) *F is level-k implementable if and only if it is Bayesian implementable*

(ii) *F is robust level-k implementable if and only if it is ex post implementable.*

PROOF:

Part (i): ($\Rightarrow$)Suppose that the social choice rule F is level-k implementable. Then there exists some mechanism $\langle \Theta, g \rangle$ and a function $m_i : \Theta_i \times \{0, \ldots, \bar{k}\} \rightarrow \Theta_i$ for each $i \in I$ such that $m = m_1 \times \cdots \times m_N$ is a level-k solution and level-k achieves F.

Suppose that there exists some player $i \in I$, and two types for player i with $\theta_i \neq \theta_i' \in \Theta_i$ such that $m_i(\theta_i, k) = m_i(\theta_i', j)$ for some $j, k \in \{1, \ldots, \bar{k}\}$. Then, it must follow that $g(m_i(\theta_i, k), m_{-i}(\theta_{-i}, 1)) \in F(\theta_i, \theta_{-i}) \cap F(\theta_i', \theta_{-i})$ for every $\theta_{-i} \in \Theta_{-i}$. This contradicts assumption A^*.

Therefore it must be true that $m_i(\theta_i, k) = m_i(\theta_i, j)$ for all $j, k \in \{1, \ldots, \bar{k}\}$.

Define $\tilde{g} : \Theta \rightarrow Y$ by $\tilde{g}(\theta) = g(m((\theta, 2)))$ for all $\theta \in \Theta$.

Then, for any $i \in I$, $\theta \in \Theta$, and $\theta' \in \Theta_i$

$$
\begin{aligned}
\sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot u_i(\tilde{g}(\theta),\theta) - \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot u_i(\tilde{g}(\theta',\theta_{-i}),\theta) \\
= \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot u_i(g(m_i(\theta_i,2),m_{-i}(\theta_{-i},2)),\theta) \\
- \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot u_i(m_i(\theta',2),m_{-i}(\theta_{-i},2)),\theta) \\
= \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot u_i(g(m_i(\theta_i,2),m_{-i}(\theta_{-i},1)),\theta) \\
- \sum_{\theta_{-i}\in\Theta_{-i}} \rho(\theta_{-i}|\theta_i)\cdot u_i(m_i(\theta',2),m_{-i}(\theta_{-i},1)),\theta) \\
\geq 0
\end{aligned}
$$

The inequality follows from the fact m is a level-k solution. Thus, the Bayesian incentive constraints hold. Further, $\langle\Theta,\tilde{g}\rangle$ achieves F because $\tilde{g}(\theta) = g(m((\theta,2))) \in F(\theta)$. Therefore, it follows by definition that F is Bayesian implementable using the mechanism$\langle\Theta,\tilde{g}\rangle$.

($\Leftarrow$) If F is Bayesian implementable there exists a mechanism $\langle\Theta, f\rangle$ that Bayesian implements F.

Set $f^i = f$ for all $i \in I$. The conditions of Proposition 2 are satisfied by the definition of Bayesian implementation. Therefore, F is level-k implementable with the restricted mechanism.

Part (ii): ($\Rightarrow$)Suppose that the social choice rule F is robust level-k implementable. Then there exists some mechanism $\langle\Theta, g\rangle$ that LDoR implements F for any LDoR type space.

Consider the LDoR type space $\mathcal{L}^*$ defined in the proof of Proposition 3.

Since, F is robust level-k implementable, there exists some mechanism $\langle\Theta, g\rangle$ and function $m_i : T_i \to \Theta_i$ for each $i \in I$ such that $m = m_1\times\cdots\times m_N$ is a LDoR solution and implements F.

Suppose that there exists some player $i \in I$, and two types for player i, t_i and t_i' with $\theta_i = \theta_i(t_i) \neq \theta_i(t_i') = \theta_i'$ such that $m_i(t_i) = m_i(t_i')$. Then, it follows that $g(m_i(t_i), m_{-i}(t_{-i})) \in F(\theta_i, \theta_{-i}(t_{-i})) \cap F(\theta_i', \theta_{-i}(t_{-i}))$ for every $t_{-i} \in T_{-i}$. But this must then mean that for every $\theta_{-i} \in \Theta_{-i}$ there exists some $z \in Y$ such that $z \in F(\theta_i, \theta_{-i}) \cap F(\theta_i', \theta_{-i})$. This contradicts assumption A^*.

Therefore it must be true that $m_i(t_i) = m_i(t_i')$ whenever $\theta_i(t_i) = \theta_i(t_i')$.

Define $\tilde{g} : \Theta \to Y$ by $\tilde{g}(\theta) = g(m(t^{2,\theta}))$ for all $\theta \in \Theta$.

Then, for any $i \in I$, $\theta \in \Theta$, and $\theta' \in \Theta_i$

$$\begin{aligned} & u_i(\tilde{g}(\theta), \theta_i) - u_i(\tilde{g}(\theta', \theta_{-i}), \theta_i) \\ & \quad = u_i(g(m(t^{2,\theta})), \theta_i) - u_i(g(m_i(t_i^{2,(\theta',\theta_{-i})}), m_{-i}(t_{-i}^{2,\theta})), \theta_i) \\ & \quad = u_i(g(m_i(t_i^{2,\theta}), m_{-i}(t_{-i}^{1,\theta})), \theta_i) - u_i(g(m_i(t_i^{2,(\theta',\theta_{-i})}), m_{-i}(t_{-i}^{1,\theta})), \theta_i) \\ & \quad \geq 0 \end{aligned}$$

The inequality follows from the fact m is a LDoR solution under $\langle \Theta, g \rangle$. Thus, the ex post incentive constraints hold. Further, $\langle \Theta, \tilde{g} \rangle$ achieves F because $\tilde{g}(\theta) = g(m(t^{2,\theta})) \in F(\theta)$. Therefore, it follows that F is ex post implementable using the mechanism$\langle \Theta, \tilde{g} \rangle$.

($\Leftarrow$) If F is ex post implementable there exists a mechanism $\langle \Theta, f \rangle$ that implements it

Set $f^i = f$ for all $i \in I$. The conditions of Proposition 4 are satisfied by definition of ex post implementable. Therefore, F is robust level-k implementable with the restricted mechanism.

□

Recall that the necessary conditions stated in Propositions 1 and 3 hold in more general environments than independent private value environments. Further, the necessary conditions do not depend on the assumptions made about level 0 behavior. This means that in these special environments, the first direction of the proofs hold independent of level 0 behavior and in general environments. In other words, in these two special environments, level-k implementation implies Bayesian implementation and robust level-k implementation implies ex post implementation regardless of the assumptions on level 0 behavior and even if the environments are not ones of independent private values.

In the next section, we focus on a specific application - bilateral trade - where we apply the results from the previous three sections. We show that in the bilateral trade environment, the social planner can always implement ex post efficient trade under robust level-k implementation. The same cannot be said for Bayesian or ex post implementation. This illustrates that the implications from Corollary 1 and Proposition 5 do not extend generally. This also illustrates that the conclusions from de Clippel et al. (ming) and Crawford (2016) regarding the limitations of level-k implementation do not extend generally.

6 Bilateral trade

6.1 Bilateral trade environment

The remainder of this paper focuses on the bilateral trade environment. Buyer's values are given by a finite set V. Seller's costs are given by a finite set$C$. The set of outcomes is given by $Y = \mathbb{R} \cup \{\emptyset\}$, where outcome $\emptyset$ indicates that the good is not traded and outcome $p \in \mathbb{R}$ indicates that the good is traded at price p. Agents have quasi-linear utility functions, $u_b : Y \times V \to \mathbb{R}$ and $u_c : Y \times C \to \mathbb{R}$. For any outcome p, the utility of a buyer with a valuation v is

$$u_b(p, v) \quad = \quad \begin{cases} v - p & \text{if } p \in \mathbb{R} \\ 0 & \text{otherwise} \end{cases}$$

and the utility of a seller with a cost c is

$$u_s(p, c) = \begin{cases} p - c & \text{if } p \in \mathbb{R} \\ 0 & \text{otherwise} \end{cases}.$$

We are interested in mechanisms that satisfy the ex post efficient social choice rule $F^*(v, c) = \{y|y \in \mathbb{R} \text{ if } v \geq c \text{ and } y = \emptyset \text{ otherwise}\}$. The ex post efficient choice rule requires trade whenever the buyer's value is above the seller's cost. We are also interested in the mechanism satisfying two additional properties: budget balance (the price paid by the buyer equals the price received by the seller) and ex post individual rationality (both the buyer and seller prefer to participate in the trading institution than receive the utility of 0). To simplify discussion, when we ask whether ex post efficiency is implementable, we do so under the condition that budget balance and ex post individual rationality are also satisfied.

6.2 Ex post efficient trade

This section contains the main result: the ex post efficient social choice rule is both level-k implementable and robust level-k implementable.

Notice that Corollary 1 does not apply in this environment. This is because the ex post efficient choice rule, F^*, is a multi-valued choice rule. If the buyer's value is above the seller's cost, ex post efficiency requires trade, but the planner does not care at what price the good is traded. This means we can potentially define two different functions, f^b and f^s, that satisfy the sufficient conditions from Propositions 2 and 4.

Also, notice that if we impose an additional assumption, A1, on the decision environment then the conditions for Proposition 5 are satisfied:

A1: For any $v, v' \in V$ with $v' < v$, there exists a $c \in C$ such that $v' \leq c \leq v$. And, for any $c, c' \in C$ with $c < c'$, there exists a $v \in V$ such that $c \leq v \leq c'$.

This means that in order to implement ex post efficient trade under level-k or robust level-k implementation, we will need to look at augmented mechanisms that increase the message space beyond that of the set of payoff types. Proposition 6 establishes the result.

Proposition 6. *The ex post efficient social choice rule, F^*, is robust level-k implementable under a mechanism that satisfies budget balance and ex post individual rationality.*

PROOF:

Define

$$f^b(v,c) = \begin{cases} c & \text{if } c \leq v \\ \emptyset & \text{otherwise} \end{cases}$$

and

$$f^s(v,c) = \begin{cases} v & \text{if } c \leq v \\ \emptyset & \text{otherwise} \end{cases}.$$

First, it is easy to see that condition (i) in Proposition 4 holds for both f^b and f^s as both assign the outcome $\emptyset$ only when $v < c$.

Now consider the utility of the buyer with value v when the seller reports cost c. Consider first the comparison of the outcomes $f^b(v,c)$ to outcomes $f^b(v',c)$ for some value $v' \in V$. There are two cases to consider:

(i)$v < c$: The utility of the buyer is 0 when reporting v and reporting any other value v' either has no effect (if $v' < c$) or achieves trade (if $v' \geq c$) with a utility of $v - c \leq 0$.

(ii)$c \leq v$: The utility of buyer is $v - c \geq 0$ when reporting v and reporting any other value v' either has no effect (if $v' > c$) or achieves outcome $\emptyset$ and utility 0.

Now, consider the comparison of the outcomes $f^b(v,c)$ to outcomes $f^s(v',c)$ for some value $v' \in V$. There are two cases to consider:

(i)$v < c$: The utility of the buyer under f^b is 0 when reporting v and reporting any other value v'under f^s either has no effect (if $v' < c$) or achieves trade (if $v' \geq c$) with a utility of $v - v' \leq 0$.

(ii)$c \leq v$: The utility of buyer is $v - c \geq 0$ under f^b when reporting v and reporting any other value v' under f^s either achieves outcome $\emptyset$ (if $v' < c$) and utility 0 or achieves trade (if $v' \geq c$) with a utility of $v - v' \leq v - c$.

Thus, the buyer has (weakly) higher utility when reporting v under f^b than reporting any other value v' under f^b or f^s regardless of the cost of the seller, c. In other words, condition (ii) in Proposition 4 is satisfied for the buyer. Analogously, condition (ii) is satisfied for the seller.

The two functions f^b and f^s satisfy the sufficient conditions for robust level-k implementation. All outcomes assigned in the mechanism in the proof of Proposition 4 are determined by $f^b$and f^s - which satisfy ex post individual rationality whenever types are truthfully reporting their payoff type. Budget balance is satisfied automatically given the specification of the environment. The environment is one of private values, thus result follows from Proposition 4.

□

This bilateral trade application provides a counter-example to the results in de Clippel et al. (forthcoming) and Crawford (2016) regarding the limitations of level-k implementation.[13] Proposition 6 illustrates that the implications from Corollary 1 and Proposition 5 do not extend generally. Level-k and robust level-k implementation can be weaker implementation concepts than both Bayesian and ex post implementation.

[13] Matsuo (1989) gives sufficient conditions for ex post efficient trade to be Bayesian implementable in the finite type bilateral trade environment. He gives examples where ex post efficient trade is not Bayesian implementable.

References

Arad, A. & Rubinstein, A. (2012). The 11-20 Money Request Game: Evaluating the Upper Bound of Level-k Reasoning. *American Economic Review*, 102(7), 3561–73.

Bergemann, D. & Morris, S. (2005). Robust Mechanism Design. *Econometrica*, 73(6), 1771–1813.

Brocas, I., Carrillo, J. D., Wang, S. W., & Camerer, C. F. (2014). Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games. *Review of Economic Studies*, 81(3), 944–970.

Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics*, 119(3), 861–898.

Costa-Gomes, M. & Crawford, V. P. (2006). Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. *American Economic Review*, 96(5), 1737–1768.

Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and Behavior in Normal-Form Games: An Experimental Study. *Econometrica*, 69(5), 1193–1235.

Costa-Gomes, M. A., Crawford, V. P., & Iriberri, N. (2013). Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications. *Journal of Economic Literature*, 51, 5–62.

Crawford, V., Kugler, Neeman, & Pauzner (2009). Behaviorally optimal auction design: examples and observations. *Journal of the European Economic Association*, 7, 377–387.

Crawford, V. P. (2016). Efficient mechanisms for level-k bilateral trading.

Crawford, V. P. & Iriberri, N. (2007). Fatal Attraction: Salience, Naivete, and Sophistication in Experimental Hide-and-Seek Games. *American Economic Review*, 97(5), 1731–1750.

de Clippel, G. (2014). Behavioral implementation. *American Economic Review*, 104(10), 2975–3002.

de Clippel, G., Saran, R., & Serrano, R. (forthcoming). Level-k mechanism design. *Review of Economic Studies*.

Eliaz, K. (2002). Fault Tolerance Implementation. *Review of Economic Studies*, 69, 589–610.

Eliaz, K. & Spiegler, R. (2006). Contracting with diversely naive agents. *Review of Economic Studies*, 73(3), 689–714.

Eliaz, K. & Spiegler, R. (2007). A mechanism design approach to speculative trade. *Econometrica*, 75(3), 875–884.

Eliaz, K. & Spiegler, R. (2008). Optimal speculative trade among large traders. *Review of Economic Design*, 12(1), 45–74.

Glazer, J. & Rubinstein, A. (1998). Motives and implementation: on the design of mechanisms to elicit opinions. *Journal of Economic Theory*, 79, 157–173.

Glazer, J. & Rubinstein, A. (2012). A model of persuasion with a boundedly rational agent. *Journal of Political Economy*, 120(6), 1057–1082.

Gorelkina, O. (2015). The expected externality mechanism in a level-k environment. *International Journal of Game Theory*, 47, 103–131.

Matsuo, T. (1989). On incentive compatible, individually rational, and ex post efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 48, 189–194.

Myerson, R. B. & Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29, 265–281.

Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *American Economic Review*, 85(5), 1313–1326.

Saran, R. (2011). Bilateral trading with naive traders. *Games and Economic Behavior*, 72, 544–557.

Severinov, S. & Deneckere, R. (2006). Screening when some agents are non-strategic: does a monopoly need to exclude? *RAND Journal of Economics*, 37, 816–840.

Stahl, D. O. & Wilson, P. W. (1994). Experimental Evidence on Player's Models of Other Players. *Journal of Economic Behavior and Organization*, 25(3), 309–327.

Stahl, D. O. & Wilson, P. W. (1995). On Player's Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1), 218–254.

Strzalecki, T. (2014). Depth of Reasoning and Higher Order Beliefs. *Journal of Economic Behavior and Organization*, 108, 108–122.

Wolitzky, A. (2016). Mechanism design with maxmin agents: Theory and an application to bilateral trade. *Theoretical Economics*, 11(3), 971–1004.

Appendix A Omitted proofs

Proof of Proposition 4:

Choose some $\bar{\theta}_i \in \Theta_i$ for each $i \in I$.

Consider the following mechanism where the message space for agent i is equal to $M_i = \Theta_i \times \{0, 1, \dots \bar{k}\} \times [-1, 1]$ and consists a report of her payoff type $\theta_i \in \Theta_i$, level $k_i \in \{0, 1, \dots, \bar{k}\}$, and a real number, $z_i \in [-1, 1]$.

Let the indicator function $I_i(z) : Z \to \{0, 1\}$ be defined as follows:

$$I_i(z) = \begin{cases} 1 & \text{if } z_i = m_j z_j \text{ for some } m_j \in \mathbb{Z}\ \forall j \in I \\ 0 & \text{otherwise} \end{cases}$$

Define $\tilde{\theta}_i : M \to \Theta_i$ and $\tilde{k}_i : M \to \{0, 1, \ldots, \bar{k}\}$ in the following way. For a given message profile $m = (\theta, k, z)$, if $I_i(z) = 1$ then $\tilde{\theta}_i(m) = \theta_i$ and $\tilde{k}_i(m) = k_i$; otherwise the planner sets $\tilde{\theta}_i(m) = \bar{\theta}_i$ and $\tilde{k}_i(m) = 0$. The planner then assigns outcomes based on the reports $\tilde{\theta} \times \tilde{k}$ according to the mechanism $g : \Theta \times \{0, \ldots, \bar{k}\}^n \to Y$ defined by

$$g(\theta \times \hat{k}) \quad = \quad \begin{cases} f^i(\theta) & \text{if } k_j < k_i - 1 \text{ for all } j \neq i \in I \\ f^1(\theta) & \text{otherwise} \end{cases}.$$

Let $\mathcal{L}^{LDoR} = \langle (T_i, k_i, \theta_i, b_i)_{i \in I}, \bar{k} \rangle$ be an LDoR type space.

Consider an agent $t_i \in T_i$ with $\theta_i(t_i) = \theta_i$ and level $k_i(t_i) = 1$. She believes that all other agent's are level 0 agents that are sending messages uniformly randomly. Thus, the probability that $I_j(z) = 1$ is zero for any $z_i \in [-1, 1]$. Hence, the type t_i agent believes that the planner will almost surely use the payoff type $\bar{\theta}_j$ for agent j and level $k_j = 0$. Further, if the type t_i agent sends a non-zero report, $z_i \neq 0$, she will expect the planner to disregard her payoff type report in favor our $\bar{\theta}_i$ with probability 1. However, if the type t_i agent sends a zero report, $z_i = 0$, she will expect the planner to use her payoff type as reported.

Thus, if she sends the message $(\theta', 1, 0)$, she will expect to receive the outcome $f^i(\theta', \bar{\theta}_{-i})$ and utility

$$u_i(f^i(\theta', \bar{\theta}_{-i}), \theta_i)$$

If she sends the message $(\theta', k_i, 0)$ she will expect to receive the outcome $f^j(\theta', \bar{\theta}_{-i})$ and utility

$$u_i(f^j(\theta', \bar{\theta}_{-i}), \theta_i)$$

for some $j \in I$.

And, if she sends the message (θ', k_i, z_i) with $z_i \neq 0$, then she will expect to receive the outcome $f^1(\bar{\theta}_i, \bar{\theta}_{-i})$, and utility

$$u_i(f^1(\bar{\theta}_i, \bar{\theta}_{-i}), \theta_i).$$

By condition (ii) we have that

$$u_i(f^i(\theta_i, \bar{\theta}_{-i}), \theta_i) \geq u_i(f^j(\theta', \bar{\theta}_{-i}), \theta_i)$$

for all $\theta' \in \Theta_i$ and all $j \in I$.

Thus, for any agent i with payoff type θ_i and level 1, reporting $(\theta_i, 1, 0)$ is a best response.

We now prove that an agent with payoff type θ_i and level k will send the message $(\theta_i, k, 0)$ by induction on the following statement: Let $k \geq 1$ and assume that if for all $l \in \{1, \ldots, k-1\}$, $\theta_j \in \Theta_j$, and $j \in I$ an agent j with payoff type θ_j and level l will report $(\theta_j, l, 0)$, then an agent i with payoff type θ_i and level k will report $(\theta_i, k, 0)$.

The result is true for $k = 1$ by the above argument. Now, consider an agent t_i with payoff type $\theta_i(t_i) = \theta_i$ and level $k_i(t_i) = k \in \{2, \ldots, \bar{k}\}$. She expects other agents that have strictly positive levels to always send reports $z_j = 0$. Thus she expects that the social planner will always take their payoff and level reports as given. She expects agents that have level 0 to randomly choose a report $z \in [-1, 1]$. Thus, she expects the planner to almost surely use payoff and level reports $\bar{\theta}_j$ and 0 for those agents.

Define the set $T_i^+ = \{t_i \in T_i | k_i(t_i) > 0\}$.

Thus, if she sends the message $(\theta_i, k, 0)$ she will expect to receive the following lottery over outcomes

$$\sum_{t_{-i} \in T_{-i}^+} \beta_i(t_{-i}|t_i) \cdot f^i(\theta_i, \theta_{-i}(t_{-i})) + \left(1 - \sum_{t_s \in T_{-i}^+} \beta_i(t_s|t_i)\right) \cdot f^i(\theta_i, \bar{\theta}_{-i}).$$

If she sends the message $(\theta_i, k_i, 0)$, $k_i \neq k$, she will expect to receive the following lottery over outcomes

$$\sum_{j \in I} \sum_{t_s \in T^+_{-i}} \beta^j_s \cdot f^j(\theta_i, \theta_{-i}(t_s)) + \left(1 - \sum_{t_s \in T^+_{-i}} \beta_i(t_s|t_i)\right) \cdot f^i(\theta_i, \bar{\theta}_{-i}).$$

for some for $\beta = ((\beta^j_s)_{t_s \in T_{-i}})_{j \in I}$ such that $\beta^j_s \in [0,1]$ and $\sum_{j \in I} \sum_{t_s \in T^+_{-i}} \beta^j_s = \sum_{t_s \in T_{-i}} \beta_i(t_s|t_i)$.

And, if she sends the message (θ_i, k_i, z_i) with $z_i \neq 0$, then she will expect to receive the following lottery over outcomes

$$\sum_{j \in I} \sum_{t_s \in T^+_{-i}} \beta^j_s \cdot f^j(\bar{\theta}_i, \theta_{-i}(t_s)) + \left(1 - \sum_{t_s \in T^+_{-i}} \beta_i(t_s|t_i)\right) \cdot f^1(\bar{\theta}_i, \bar{\theta}_{-i})$$

for some for $\beta = ((\beta^j_s)_{t_s \in T_{-i}})_{j \in I}$ such that $\beta^j_s \in [0,1]$ and $\sum_{j \in I} \sum_{t_s \in T_{-i}} \beta^j_s = \sum_{t_s \in T_{-i}} \beta_i(t_s|t_i))$.

By condition (ii) we have that

$$u_i(f^i(\theta), \theta_i) \geq u_i(f^j(\theta', \theta_{-i}), \theta_i)$$

for all $\theta' \in \Theta_i$ and all $j \in I$.

It must also then also be true that

$$\sum_{t_s \in T^+_{-i}} \beta_i(t_s|t_i) \cdot f^i(\theta_i, \theta_{-i}(t_s)) + \left(1 - \sum_{t_s \in T_{-i}} \beta_i(t_s|t_i)\right) \cdot f^i(\theta_i, \bar{\theta}_{-i})$$
$$\geq \sum_{j \in I} \sum_{t_s \in T^+_{-i}} \beta^j_s \cdot f^j(\theta', \theta_{-i}(t_{-i})) + \left(1 - \sum_{t_s \in T_{-i}} \beta_i(t_s|t_i)\right) \cdot f^l(\theta', \bar{\theta}_{-i}).$$

for any $\beta = ((\beta^j_s)_{t_s \in T_{-i}})_{j \in I}$ such that $\beta^j_s \in [0,1]$ and $\sum_{j \in I} \sum_{t_s \in T^+_{-i}} \beta^j_s = \sum_{t_s \in T_{-i}} \beta_i(t_s|t_i))$ and for all $\theta' \in \Theta_i$ and $l \in I$.

Thus, for agent i with payoff type θ_i and level k, reporting $(\theta_i, k, 0)$ is a best response.

Therefore, if we define $m_i(t_i) = (\theta_i(t_i), k_i(t_i), 0)$ for all $t_i \in T_i$ with $k_i(t_i) \in \{1, \ldots, \bar{k}\}$, then m is a level-k solution and m achieves F by condition (i).

Since the above holds for an arbitrary LDoR type space, F is robust level-k implementable.

$\square$