KELLEY-MORSE SET THEORY DOES NOT PROVE THE CLASS FODOR PRINCIPLE

VICTORIA GITMAN, JOEL DAVID HAMKINS, AND ASAF KARAGILA

Abstract. We show that Kelley-Morse set theory does not prove the class Fodor principle, the assertion that every regressive class function $F : S \to \mathrm{Ord}$ defined on a stationary class S is constant on a stationary subclass. Indeed, it is relatively consistent with KM for any infinite λ with $\omega \leq \lambda \leq \mathrm{Ord}$ that there is a class function $F : \mathrm{Ord} \to \lambda$ that is not constant on any stationary class. Strikingly, it is consistent with KM that there is a class $A \subseteq \omega \times \mathrm{Ord}$, such that each section $A_n = \{\, \alpha \mid (n, \alpha) \in A \,\}$ contains a class club, but $\bigcap_n A_n$ is empty. Consequently, it is relatively consistent with KM that the class club filter is not σ-closed.

1. Introduction

The *class Fodor principle* is the assertion that every regressive class function $F : S \to \mathrm{Ord}$ defined on a stationary class S is constant on a stationary subclass of S. This statement can be expressed in the usual second-order language of set theory, and the principle can therefore be sensibly considered in the context of any of the various second-order set-theoretic systems, such as Gödel-Bernays GBC set theory or Kelley-Morse KM set theory. Just as with the classical Fodor's lemma in first-order set theory, the class Fodor principle is equivalent, over a weak base theory, to the assertion that the class club filter is normal (see theorem 5). We shall investigate the strength of the class Fodor principle and try to find its place within the natural hierarchy of second-order set theories. We shall also define and study weaker versions of the class Fodor principle.

If one tries to prove the class Fodor principle by adapting one of the classical proofs of the first-order Fodor's lemma, then one inevitably finds oneself needing to appeal to a certain second-order class-choice principle, which goes beyond the axiom of choice and the global choice principle, but which is not available in Kelley-Morse set theory [GH]. For example, in one standard proof, we would want for a given Ord-indexed sequence of non-stationary classes to be able to choose for each member of it a class club that it misses. This would be an instance of class-choice, since we seek to choose classes here, rather than sets. The class choice principle $\mathrm{CC}(\Pi^0_1)$, it turns out, is sufficient for us to make these choices, for this principle states that if every ordinal α admits a class A witnessing a Π^0_1-assertion $\varphi(\alpha, A)$, allowing class parameters, then there is a single class $B \subseteq \mathrm{Ord} \times V$, whose slices B_α witness $\varphi(\alpha, B_\alpha)$; and the property of being a class club avoiding a given class is Π^0_1 expressible.

The third author was partially supported by the Austrian Science Foundation (FWF), Grant I 3081, and by the Royal Society grant no. NF170989.

Commentary about this article can be made at http://jdh.hamkins.org/km-does-not-prove-class-fodor.

Thus, the class Fodor principle, and consequently also the normality of the class club filter, is provable in the relatively weak second-order set theory $\mathrm{GBC}+\mathrm{CC}(\Pi^0_1)$, as shown in theorem 7. This theory is known to be weaker in consistency strength than the theory $\mathrm{GBC}+\Pi^1_1$-comprehension, which is itself strictly weaker in consistency strength than KM.

But meanwhile, although the class choice principle is weak in consistency strength, it is not actually provable in KM; indeed, even the weak fragment $\mathrm{CC}(\Pi^0_1)$ is not provable in KM. Those results were proved several years ago by the first two authors [GH], but they can now be seen as consequences of the main result of this article (see corollary 15). In light of that result, however, one should perhaps not have expected to be able to prove the class Fodor principle in KM.

Indeed, it follows similarly from arguments of the third author in [Kar18] that if κ is an inaccessible cardinal, then there is a forcing extension $V[G]$ with a symmetric submodel M such that $V^M_\kappa = V_\kappa$, which implies that $\mathcal{M} = (V_\kappa, \in, V^M_{\kappa+1})$ is a model of Kelley-Morse, and in $\mathcal{M}$, the class Fodor principle fails in a very strong sense.

In this article, adapting the ideas of [Kar18] to the second-order set-theoretic context and using similar methods as in [GH], we shall prove that every model of KM has an extension in which the class Fodor principle fails in that strong sense: there can be a class function $F : \mathrm{Ord} \to \omega$, which is not constant on any stationary class. In particular, in these models, the class club filter is not σ-closed: there is a class $B \subseteq \omega \times \mathrm{Ord}$, each of whose vertical slices B_n contains a class club, but $\bigcap B_n$ is empty.

Main Theorem. *Kelley-Morse set theory* KM*, if consistent, does not prove the class Fodor principle. Indeed, if there is a model of* KM*, then there is a model of* KM *with a class function $F : \mathrm{Ord} \to \omega$, which is not constant on any stationary class; in this model, therefore, the class club filter is not σ-closed.*

We shall also investigate various weak versions of the class Fodor principle.

Definition 1.

(1) For a cardinal κ, the *class κ-Fodor principle* asserts that every class function $F : S \to \kappa$ defined on a stationary class $S \subseteq \mathrm{Ord}$ is constant on a stationary subclass of S.
(2) The *class* $<$Ord-*Fodor principle* is the assertion that the κ-class Fodor principle holds for every cardinal κ.
(3) The *bounded class Fodor principle* asserts that every regressive class function $F : S \to \mathrm{Ord}$ on a stationary class $S \subseteq \mathrm{Ord}$ is bounded on a stationary subclass of S.
(4) The *very weak class Fodor principle* asserts that every regressive class function $F : S \to \mathrm{Ord}$ on a stationary class $S \subseteq \mathrm{Ord}$ is constant on an unbounded subclass of S.

We shall separate these principles as follows.

Theorem 2. *Suppose* KM *is consistent.*

(1) *There is a model of* KM *in which the class Fodor principle fails, but the class* $<$Ord-*Fodor principle holds.*
(2) *There is a model of* KM *in which the class ω-Fodor principle fails, but the bounded class Fodor principle holds.*

(3) *There is a model of* KM *in which the class ω-Fodor principle holds, but the bounded class Fodor principle fails.*
(4) GB^- *proves the very weak class Fodor principle.*

Finally, we show that the class Fodor principle can neither be created nor destroyed by set forcing.

Theorem 3. *The class Fodor principle is invariant by set forcing over models of* GBC^-. *That is, it holds in an extension if and only if it holds in the ground model.*

Let us conclude the introduction by mentioning the following easy negative instance of the class Fodor principle for certain GBC models. This argument seems to be a part of set-theoretic folklore. Namely, consider an ω-standard model of GBC set theory M having no V_κ^M that is a model of ZFC. A minimal transitive model of ZFC, for example, has this property. Inside M, let $F(\kappa)$ be the least n such that V_κ^M fails to satisfy Σ_n-collection. This is a definable class function $F : \mathrm{Ord}^M \to \omega$ in M, but it cannot be constant on any stationary class in M, because by the reflection theorem there is a class club of cardinals κ such that V_κ^M satisfies Σ_n-collection.

2. Second-order set theories

We shall express all our second-order set-theories in a two-sorted language, with separate variables for sets and classes, and a binary relation $\in$ used both for the membership relation on sets $a \in b$ and also for membership of a set in a class $a \in A$. Models of these theories can be taken to have the form $\mathcal{M} = \langle M, \in^{\mathcal{M}}, \mathscr{C} \rangle$, where M is the collection of sets in the model and $\mathscr{C}$ is the collection of classes A with $A \subseteq M$.

Let us briefly review the principal second-order set theories we shall consider. Gödel-Bernays set theory GB asserts ZF for the first-order part of the structure, plus extensionality for classes, the replacement axiom for class functions, and the class-comprehension axiom, which asserts that every first-order formula, allowing set and class parameters, defines a class. The theory $\mathrm{GB} + \mathrm{AC}$ includes the axiom of choice for sets, and the theory GBC includes the global choice principle, which asserts that there is a well-order of the sets in type Ord. The natural weakenings GB^- and GBC^- drop the power set axiom, asserting only ZF^- for the first-order part (when dropping the power set axiom, be sure to use the collection and separation schemes in ZF^-, instead of just replacement, in light of [GHJ16]).

A hierarchy of theories grows above GBC by strengthening the class comprehension axiom into the second-order language. The theory $\mathrm{GBC} + \Pi^1_n$-comprehension, for example, allows class comprehension for Π^1_n-expressible properties, with class parameters. Growing in power as n increases, this hierarchy culminates ultimately in Kelley-Morse set theory KM, which is GBC together with the full second-order class comprehension axiom.

Recent work has revealed a central role in the hierarchy for the principle of elementary transfinite recursion ETR, which asserts that every first-order definable class recursion along a class well-order, including orders longer than Ord, has a solution (see [GH16, GHH$^+$17]). The principle is naturally stratified according to the length of the class well order. The principle $\mathrm{ETR}_{\mathrm{Ord}}$, for example, asserting solutions for class recursions of length Ord, is equivalent to the assertion that every class forcing notion admits a forcing relation [GHH$^+$17]; and the principle ETR_ω, for class recursions of length ω, proves that every structure $\langle V, \in, A\rangle$, for any class A, admits a first-order truth predicate; the Tarskian recursive definition of truth is, after all, a recursion of length ω on formulas. In this way, the restricted ETR principles span the region between GBC and GBC + Π^1_1-comprehension, which proves ETR.

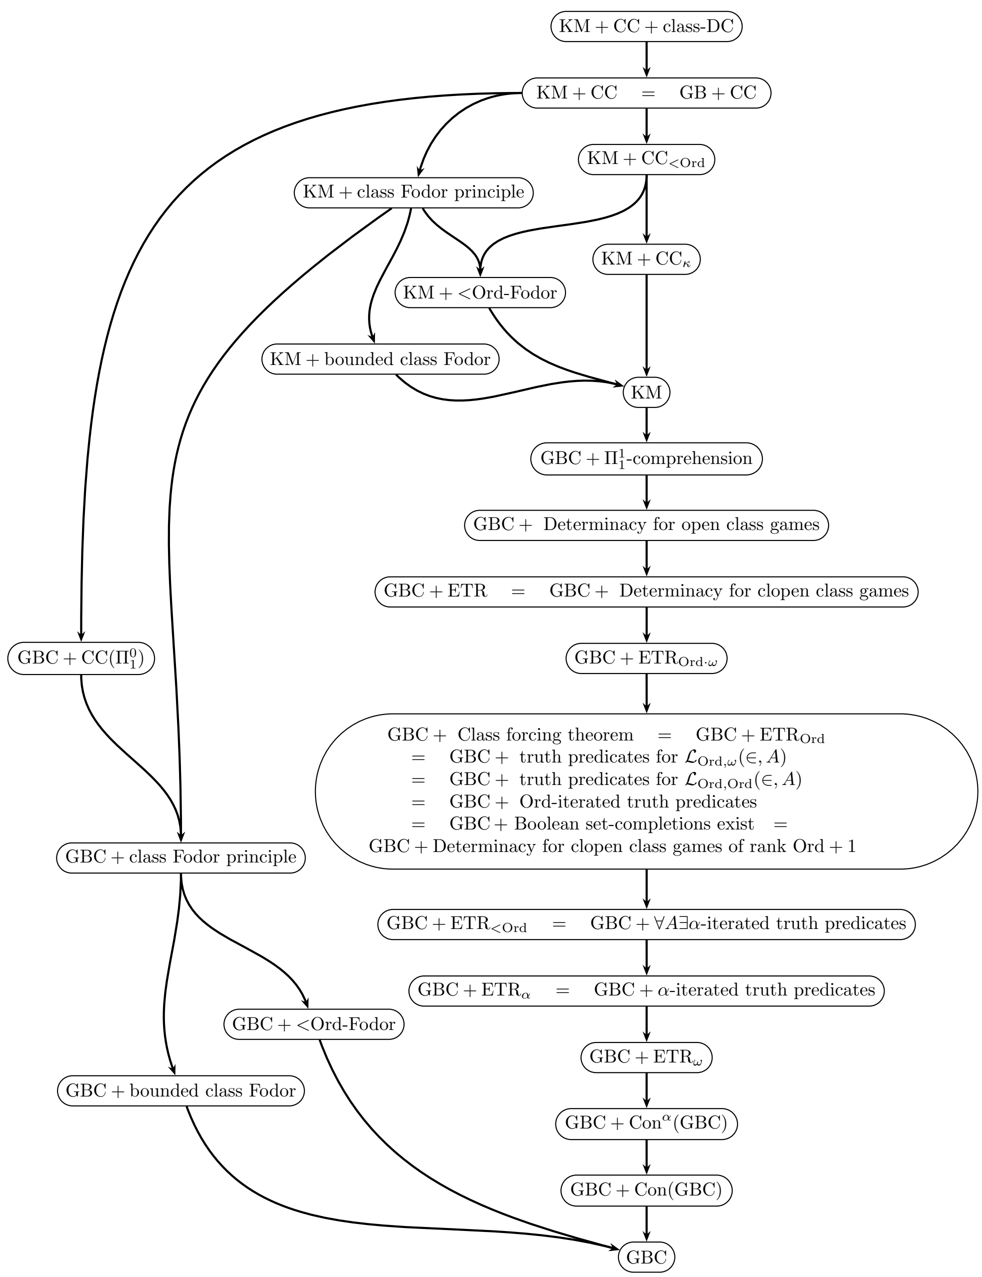

FIGURE 1. The hierarchy of second-order set theories

Figure 1 shows the implication diagram for the hierarchy of second-order theories (except that there should be additional arrows, which we omitted for clarity, from KM + bounded class Fodor and KM + <Ord-Fodor to the corresponding respective GBC theories at the bottom left). Many of the arrows in the diagram, but not all, are also strict consistency strength implications. For example, the ETR principles rise in consistency strength. Yet, we've mentioned that KM + CC is equiconsistent with KM; it follows that all the various class Fodor principles are equiconsistent over KM. We don't yet know whether the corresponding fact holds over GBC, and in section 10, we ask several questions about the strength of the class Fodor principles over GBC, where we do not know the exact consistency strength. In terms of consistency strength, to be sure, the entire diagram does not vary much, since even KM + CC + DC at the top is strictly below ZFC+ one inaccessible cardinal.

An enumerated collection of classes $\langle X_a \mid a \in A\rangle$ indexed by a class A in a model of second-order set theory is *coded* in the model, if there is a class $X \subseteq A \times V$, whose slices $X_a = \{\, x \mid (a, x) \in X \,\}$ are exactly the corresponding classes on the sequence.

The *class choice* principle CC is the assertion that for any class A and class parameter Z, if for every $a \in A$ there is a class X with $\varphi(a, X, Z)$, where φ is any assertion in the second-order language of set theory, then there is a coded sequence of classes choosing such witnesses; that is, there is a class $X \subseteq A \times V$, such that $\varphi(a, X_a, Z)$ for every $a \in A$. Stratifying this axiom scheme, let $\mathrm{CC}(\Pi^0_1)$ be the principle resulting when we allow φ only of complexity Π^0_1, and similarly for other levels of complexity. The axiom CC_κ is the case $A = \kappa$.

The first and second authors showed in [GH] that Kelley-Morse set theory does not prove class choice, even for Π^0_1 assertions. Meanwhile, the theories KM and KM + CC are equiconsistent, shown by Marek and Mostowski [MM75], because the 'constructible universe' of a KM-model (where one not only restricts the sets to those in the constructible hierarchy up to Ord, but also restricts the classes to those with codes appearing in the constructible hierarchy at some coded meta-ordinal stage above Ord in the corresponding unrolled structure obtained from the model) satisfies KM + CC. This equivalence holds level-by-level, so that GBC+Π^1_n-comprehension is equiconsistent with GBC+Π^1_n-comprehension together with $\mathrm{CC}(\Pi^1_n)$. Similarly, Kameryn Williams [Wil18] proved that the consistency of GBC+ETR implies the consistency of GBC together with class choice for first-order assertions, thereby providing an upper bound on the consistency strength of GBC with the class Fodor principle; his dissertation [Wil18] is an excellent resource for all these matters. Note that because $\mathrm{GBC} + \mathrm{CC}(\Pi^0_1)$ proves the class Fodor principle, while KM does not, our main theorem provides natural instances of non-linearity in the hierarchy of second-order set theories.

We say that a class $A \subseteq \mathrm{Ord}$ is in the *class club filter*, if A contains a class club. The filter itself is not a class, nor even coded as a class; it is nevertheless sensible to refer to the filter in our models by viewing it as a second-order definable meta-class, a predicate on the classes. The class club filter is $\leq \alpha$*-closed* for an ordinal α, if whenever $\langle A_\xi \mid \xi < \alpha\rangle$ is a coded collection of classes $A_\alpha \subseteq \mathrm{Ord}$, each in the class club filter, then the intersection $\bigcap_\alpha A_\alpha$ is also in the class club filter. This is expressible in the second-order language of set theory. Similarly, the class club filter is *normal*, if whenever $\langle A_\alpha \mid \alpha \in \mathrm{Ord}\rangle$ is a coded collection of classes in the class club filter, then its diagonal intersection is also in the class club filter.

3. Positive instances of the class Fodor principle

Let us begin with the observation that GB^-, without any choice principle at all, proves the very weak class Fodor principle. A version of this theorem was observed by Walter Neumer in 1951 [Neu51].

Theorem 4 (The very weak class Fodor principle)**.** *Assume* GB^-*. Then every regressive class function* $F : S \to \mathrm{Ord}$ *defined on a stationary class* S *is constant on an unbounded subclass* $A \subseteq S$*.*

Proof. Assume that S is a stationary class of ordinals and $F(\gamma) < \gamma$ for all $\gamma \in S$. If F is not constant on any unbounded subclass of S, then for every ordinal γ, there is an ordinal bound β_γ on the pre-image of γ. That is, $F(\alpha) \neq \gamma$ for all $\alpha \geq \beta_\gamma$. The function $\gamma \mapsto \beta_\gamma$ is first-order definable from F and therefore forms a class in GB^-. Let C be the class of closure points of this function, so $\theta \in C$ if and only if $\beta_\gamma < \theta$ for every $\gamma < \theta$. This is a closed unbounded class of ordinals. Since S is stationary, there is some $\theta \in C \cap S$. But notice that $F(\theta) \neq \gamma$ for every $\gamma < \theta$, since $\beta_\gamma < \theta$ for all such γ. This contradicts our assumption that F is regressive on S. $\square$

It follows that for any regressive function $F : S \to \mathrm{Ord}$ defined on a stationary class S and any coded sequence of closed unbounded classes $\langle C_\alpha \mid \alpha \in \mathrm{Ord}\rangle$, the function F has constant value γ on a subclass $T \subseteq S$ such that $T \cap C_\alpha$ is unbounded for every α. This follows as an instance because the diagonal intersection $C = \triangle_\alpha C_\alpha$ is club and F is regressive on $S \cap C$, which is itself a stationary class, and so there is an unbounded class $T \subseteq S \cap C$ on which F is constant.

Note that one cannot omit or even weaken the assumption in the theorem that S is stationary in theorem 4, because if S is not stationary, then there is a closed unbounded class of ordinals $C \subseteq \mathrm{Ord}$ disjoint from S, and in this case the function $F(\gamma) = \sup(C \cap \gamma)$ is regressive on S, but not constant on any unbounded class.

The proof of theorem 4 can be used to show in ZF^- similarly that every regressive function on a regular cardinal has an unbounded fiber.

We observe next that the classical equivalence between Fodor's lemma and the normality of the class club filter translates to the class version.

Theorem 5. *The following are equivalent over* GB^-.

(1) *The class Fodor principle.*
(2) *The class club filter is normal.*

Proof. $(1 \to 2)$ Assume that the class Fodor principle holds, and suppose that we have a coded Ord-sequence $\langle A_\alpha \mid \alpha < \mathrm{Ord}\rangle$ of classes A_α, each of which contains a class club. Note that we cannot necessarily choose these clubs, since GB^- and, indeed, even KM does not prove the class-choice principle CC. But we can appeal to the class Fodor principle. If the diagonal intersection $A = \triangle_\alpha A_\alpha$ does not contain a class club, then the complement $S = \mathrm{Ord} - A$ is stationary. For $\beta \in S$, let $F(\beta)$ be the least $\alpha < \beta$ for which $\beta \notin A_\alpha$. So $F : S \to \mathrm{Ord}$ is a regressive class function on the stationary class S. If $F \restriction T$ has constant value α for some $T \subseteq S$, then T is disjoint from A_α, and hence disjoint from a class club, and therefore not stationary. So F is not constant on any stationary class. Therefore, the diagonal intersection $\triangle_\alpha A_\alpha$ must contain a class club after all, and so the class club filter is closed under diagonal intersection.

($2 \to 1$) Conversely, suppose that the class club filter is closed under diagonal intersection and that $F : S \to \mathrm{Ord}$ is a regressive class function defined on a stationary class $S \subseteq \mathrm{Ord}$. Let $B_\alpha = F^{-1}(\{\alpha\}) = \{\, \beta \mid F(\beta) = \alpha \,\}$. If some B_α is stationary, then F is constant on a stationary class, and we're done. Otherwise, the complements $A_\alpha = \mathrm{Ord} - B_\alpha$ each contain a class club. In this case, it follows by our assumption that the diagonal intersection $A = \triangle_\alpha A_\alpha$ also contains a class club. But note that $\beta \in A \cap S$ implies $F(\beta) \neq \alpha$ for any $\alpha < \beta$, contrary to the fact that F was regressive on S. □

Recall from the introduction that for a cardinal κ, the class κ-Fodor principle is the weakening of the class Fodor principle asserting that every class function $F : S \to \kappa$ defined on a stationary class $S \subseteq \mathrm{Ord}$ is constant on a stationary subclass of S.

Theorem 6. *The following are equivalent in* GB^- *for any ordinal* κ*.*

1. *The class* κ*-Fodor principle.*
2. *The class club filter is* $\leq \kappa$*-closed.*

The proof is analogous to the proof of theorem 5, if we everywhere replace diagonal intersections of length Ord by intersections of length κ.

Finally, as we mentioned in the introduction, we can implement one of the proofs of the classical Fodor's lemma to prove the class Fodor principle, if we assume a sufficient class-choice principle.

Theorem 7. *The theory* GB^- *augmented with the class-choice principle* $\mathrm{CC}(\Pi^0_1)$ *proves the class Fodor principle. More generally,* $\mathrm{GB}^- + \mathrm{CC}_\kappa(\Pi^0_1)$ *proves the class* κ*-Fodor principle.*

Proof. Suppose toward contradiction that $F : S \to \mathrm{Ord}$ is regressive on a stationary class $S \subseteq \mathrm{Ord}$, but not constant on any stationary subclass of S. So for each ordinal α, there is a class club C with $F(\gamma) \neq \alpha$ for all $\gamma \in C$. This is a Π^0_1 property of the class C, with class parameter F. By the class-choice principle $\mathrm{CC}(\Pi^0_1)$, therefore, we may choose such clubs, and so there is a coded sequence of classes $\langle C_\alpha \mid \alpha \in \mathrm{Ord}\rangle$, where $\gamma \in C_\alpha \to F(\gamma) \neq \alpha$. Let $C = \triangle_{\alpha\in\mathrm{Ord}} C_\alpha$ be the diagonal intersection of these class clubs, which is itself a class club and therefore meets S. So there is some $\gamma \in C \cap S$. But since $F(\gamma) \neq \alpha$ for all $\alpha < \gamma$, this contradicts our assumption that F was regressive on S.

The same argument works for $F : S \to \kappa$, using only $\mathrm{CC}_\kappa(\Pi^0_1)$, since we need only pick κ many clubs in this case. □

We will show in Section 7 that it is consistent relative to KM that the class <Ord-Fodor principle holds, but the class Fodor principle fails.

Recall from the introduction that the bounded class Fodor principle is the assertion that every regressive class function $F : S \to \mathrm{Ord}$ on a stationary class $S \subseteq \mathrm{Ord}$ is bounded on a stationary subclass $T \subseteq S$.

Lemma 8. *The class Fodor principle is equivalent over* GB^- *to the conjunction of the bounded class Fodor principle and the class* <Ord*-Fodor principle.*

Proof. The class Fodor principle clearly implies both the bounded class Fodor principle and the class <Ord-Fodor principle. Conversely, if $F : S \to \mathrm{Ord}$ is regressive on a stationary class, then by the bounded class Fodor principle, we can first restrict to a stationary class $T \subseteq S$ on which F is bounded by some ordinal β, and

then by the β-Fodor principle, we can further restrict to a stationary subclass on which F is constant, as desired. $\square$

Consequently, once the bounded class Fodor principle holds, then the full class Fodor principle is equivalent to the class $<$Ord-Fodor principle. It follows from the theorems in section 7 that neither the bounded class Fodor principle nor the class $<$Ord-Fodor principle implies the other.

4. Class forcing in second-order set theories

Our main argument will involve class forcing over models of second-order set theory, and so let us briefly review some of the basics.

A class forcing notion $\mathbb{P}$ is simply a class partial pre-order, with maximal element $\mathbb{1}$, and a $\mathbb{P}$-*name* is a set or class consisting of pairs (τ, p), where τ is a $\mathbb{P}$-name set and $p \in \mathbb{P}$. The check names $\check{x} = \{\langle \check{y}, \mathbb{1}\rangle \mid y \in x\}$, for example, are defined recursively, and similarly for classes $\check{A} = \{\langle \check{a}, \mathbb{1}\rangle \mid a \in A\}$.

A class forcing notion $\mathbb{P}$ admits a *forcing relation*, if there is a relation $\Vdash_{\mathbb{P}}$ satisfying the *forcing relation recursion*, as described in [GHH$^+$17]. This is a class recursion on $\mathbb{P}$ names, by which the forcing relation $p \Vdash_{\mathbb{P}} \sigma = \tau$ and $p \Vdash_{\mathbb{P}} \sigma \in \tau$ on atomic formulas is seen to obey the expected recursive properties. This recursion is expressible internally in the language of second-order set theory, making no reference to countable models, generic filters or the truth and definability lemmas. One easily extends the forcing relation from atomic formulas, when it exists, to forcing relations on other formulas—the atomic forcing relation is the key difficult case.

The theory GBC, it turns out, is insufficient to prove that every class forcing notion admits a forcing relation, although the stronger theory GBC+ETR can prove this, and indeed, the main theorem of [GHH$^+$17] is that the existence of forcing relations for all class forcing is equivalent over GBC to the principle $\mathrm{ETR}_{\mathrm{Ord}}$.

Meanwhile, GBC is able to prove that forcing relations exist for certain particularly well-behaved forcing notions, including all the forcing we shall employ in this article. For instance, in GB^- every pre-tame forcing notion, and consequently also every tame forcing notion (see [Fri00] for definitions), admits a definable forcing relation. Such forcing preserves both KM and $\mathrm{KM} + \mathrm{CC}$ to the forcing extension.

Theorem 9 (Stanley, see [HKS18])**.** *Assume* GB^-*. Every pre-tame class forcing notion* $\mathbb{P}$ *admits a first-order* $\mathbb{P}$*-definable forcing relation for atomic formulas (and hence forcing relations for any particular formula).*

Theorem 10 ([Ant18])**.** *Tame forcing preserves* KM *to forcing extensions.*

Theorem 11 ([AFG])**.** *Tame forcing preserves* $\mathrm{KM} + \mathrm{CC}$ *to forcing extensions.*

Strengthening tameness, a class forcing notion $\mathbb{P}$ is $<$Ord-*distributive*, if for any ordinal β and any coded β-sequence $\langle D_\alpha \mid \alpha < \beta\rangle$ of open dense classes $D_\alpha \subseteq \mathbb{P}$, the intersection $\bigcap_{\alpha<\beta} D_\alpha$ is dense in $\mathbb{P}$. It is an immediate consequence of $<$Ord-distributivity that such forcing is pre-tame, and consequently, such forcing admits a definable forcing relation. The standard distributivity arguments then show that such forcing adds no new sets, and so it is also tame.

Strengthening further, we note that if a forcing notion has κ-closed dense subclasses for every ordinal κ, then it follows easily that it is $<$Ord-distributive, and consequently it admits a definable forcing relation and adds no new sets.

Typically, the forcing language for forcing over models of second-order set theory includes the membership relation $\in$ and constants for every $\mathbb{P}$-name, for both sets and classes. The case of $\check{V}$ amounts to a predicate for the ground model sets, where $p \Vdash \sigma \in \check{V}$ just in case it is dense below p that a condition forces $\sigma = \check{x}$ for some set x in V.

In this article, we should like to augment the forcing language with a predicate $\check{\mathscr{C}}$ also for the ground-model classes, defining for a class name $\dot{X}$ that $p \Vdash \dot{X} \in \check{\mathscr{C}}$ just in case it is dense below p to force $\dot{X} = \check{Y}$ for some class $Y \in \mathscr{C}$. Note that this definition involves second-order quantifiers, which will affect the complexity of forced second-order assertions making use of the $\check{\mathscr{C}}$ class predicate. But in models of KM, as in our application, this will cause no problem.

In the proof of our main theorem, we shall make a certain homogeneity argument concerning our main forcing notion. Let us therefore define here that a forcing notion $\mathbb{P}$ is *locally homogeneous*, if for every pair of conditions p and p', there are stronger conditions $q \leq p$ and $q' \leq p'$ such that $\mathbb{P} \upharpoonright q \cong \mathbb{P} \upharpoonright q'$. By carrying the automorphism through the forcing relation, just as in the familiar case of set forcing, it follows from this that if a condition p forces $\varphi(\check{a}, \check{A})$, a statement involving only check names, then every condition $q \in \mathbb{P}$ forces $\varphi(\check{a}, \check{A})$.

Finally, we describe how to construct the forcing extension models. Suppose that $\mathbb{P}$ is a class forcing notion in the model of second-order set theory $\mathcal{W} = \langle W, \in^{\mathcal{W}}, \mathscr{C} \rangle$, which has a forcing relation for it. A filter $G \subseteq \mathbb{P}$ is $\mathcal{W}$-generic, if $G \cap D$ is nonempty for every dense class $D \subseteq \mathbb{P}$ in $\mathcal{W}$. Every countable model, for example, even an ill-founded model, is easily seen to admit such generic filters.

Define the equivalence relation $\sigma =_G \tau$ to hold when there is some $p \in G$ with $p \Vdash \sigma = \tau$; and the membership relation $\sigma \in_G \tau$, when some $p \in G$ has $p \Vdash \sigma \in \tau$. Similarly define the class relations $\dot{X} =_G \dot{Y}$ and $\sigma \in_G \dot{X}$. The forcing extension $\mathcal{W}[G] = \langle W[G], \in^{\mathcal{W}[G]}, \mathscr{C}[G] \rangle$ is then formed as equivalence classes of set and class names by the relation $=_G$, which is a congruence with respect to $\in_G$. When the model is transitive or at least well-founded, one may equivalently view this construction as proceeding by a recursive interpretation-of-names definition of values $\sigma_G = \mathrm{val}(\sigma, G) = \{\, \mathrm{val}(\tau, G) \mid \exists p \in G\ \langle \tau, p \rangle \in \sigma \,\}$, but in the general case, that recursion is not actually sensible. One can augment the structure with the ground-model predicates $\check{V}$ and $\check{\mathscr{C}}$, as we discussed above.

Finally, one proves what amounts to a forcing analogue of the Łoś-theorem, namely, $\mathcal{W}[G] \models \varphi$ just in case there is some $p \in G$ with $p \Vdash \varphi$, fulfilling the desired alignment between what is forced and what is true in the forcing extension. See further details in [GHH$^+$17, §3].

5. The class-club-shooting forcing

We shall now introduce and study the class club-shooting forcing, which enables one to shoot a class club through any fat-stationary class of ordinals. Generalizing the notion from sets, a class $S \subseteq \mathrm{Ord}$ is *fat stationary*, if for every class club $C \subseteq \mathrm{Ord}$ and ordinal β, there is a closed copy of $\beta + 1$ inside $S \cap C$. The *class club-shooting* forcing $\mathbb{Q}_S$, for a fat-stationary class S, is the partial order of all closed sets $c \subseteq S$, ordered by end-extension. Let us now establish the basic properties of this forcing.

Lemma 12. *Assume* GBC, *and suppose that* $S \subseteq \mathrm{Ord}$ *is a fat-stationary class of ordinals. Then the class-club-shooting forcing* $\mathbb{Q}_S$ *has the following properties:*

(1) $\mathbb{Q}_S$ *is* $<$Ord-*distributive;*
(2) *it is therefore tame, admits a forcing relation and does not add sets;*
(3) *it adds a class club* $C \subseteq S$*;*
(4) *it preserves the fat-stationarity of any fat stationary subclass* $T \subseteq S$*; and*
(5) *it is locally homogeneous.*

Proof. Assume GBC, and suppose that $S \subseteq \mathrm{Ord}$ is a fat-stationary class of ordinals. Let $\mathbb{Q}_S$ be the class-club-shooting forcing for S. The fat stationarity of S ensures that any condition in $\mathbb{Q}_S$ can be extended to a longer condition, with order type as long as desired, and so the union of the generic filter will be a class club $C \subseteq S$.

It is easy to see that $\mathbb{Q}_S$ is locally homogeneous: if c and d are conditions in $\mathbb{Q}_S$, then extend them to conditions c^+ and d^+ having the same supremum, and consider the map π defined on extensions of c^+ that simply replaces the c^+-initial segment of a condition with d^+. This is an isomorphism of $\mathbb{Q}_S \upharpoonright c^+$ with $\mathbb{Q}_S \upharpoonright d^+$, verifying local homogeneity.

To prove that the forcing is $<$Ord-distributive, consider any condition $c_0 \in \mathbb{Q}_S$ and any coded β-sequence $\vec{D} = \langle D_\alpha \mid \alpha < \beta\rangle$ of open dense classes $D_\alpha \subseteq \mathbb{Q}_S$. We seek an extension of c_0 that is simultaneously in all D_α. To find such an extension, apply the reflection theorem to find a class club E of ordinals θ for which $\langle V_\theta, \in, S \cap \theta, \vec{D} \upharpoonright \theta\rangle \prec_{\Sigma_2} \langle V, \in, S, \vec{D}\rangle$. These structures for $\theta \in E$ therefore form a continuous Σ_2-elementary chain. Since S is fat stationary, there is a continuous $(\beta+1)$-sequence $\langle \theta_\alpha \mid \alpha \leq \beta\rangle$ inside $E \cap S$, as high as we like. We may assume $c_0 \in V_{\theta_0}$ and that $\beta < \theta_0$.

We construct a descending sequence of conditions $c_0 \geq c_1 \geq \cdots \geq c_\beta$ below c_0, ensuring that $c_\alpha \in V_{\theta_{\alpha+1}}$, that $\theta_\alpha \in c_{\alpha+1}$ and that $c_{\alpha+1} \in D_\alpha$ for all $\alpha \leq \beta$. Given c_α, add the point θ_α, which is in S and therefore allowed, and extend $c_\alpha \cup \{\theta_\alpha\}$ to a condition $c_{\alpha+1} \in D_\alpha$ inside $V_{\theta_{\alpha+1}}$, which is possible because this is Σ_2-elementary in $\langle V, \in, S, \vec{D}\rangle$. At a limit stage λ, let $c_\lambda = \cup_{\alpha<\lambda} c_\alpha \cup \{\theta_\lambda\}$, which simply takes the union of the earlier conditions and adds the supremum θ_λ. This is a legal condition because all θ_α are in S. After β steps of this construction, the condition c_β is an extension of c_0 that is inside all the open dense classes D_α, and so we have verified distributivity.

It follows that the forcing $\mathbb{Q}_S$ is tame, and therefore has a forcing relation, and from this, it follows easily from $<$Ord-distributivity that the forcing $\mathbb{Q}_S$ adds no new sets, as we explained in the previous section.

What remains is to show that the forcing preserves the fat-stationarity of any fat stationary subclass $T \subseteq S$. Suppose that $T \subseteq S$ is fat stationary. We want to see that T remains fat stationary in the forcing extension $\mathcal{V}[C]$. To see this, suppose that $\dot{D}$ is a $\mathbb{Q}_S$-name for any class club, and fix any ordinal β. Let $\Vdash$ be a class forcing relation for $\mathbb{Q}_S$ for Σ^0_2 assertions in the forcing language with parameters for T, S, and $\dot{D}$. As above, fix a continuous Σ_2-elementary $(\omega\beta+1)$-chain of structures

$$\langle V_\theta, \in, T \cap \theta, S \cap \theta, \dot{D} \cap V_\theta, \Vdash \upharpoonright V_\theta\rangle \prec_{\Sigma_2} \langle V, \in, T, S, \dot{D}, \Vdash\rangle.$$

We build a descending sequence of conditions $c_0 \geq c_1 \geq \cdots \geq c_{\omega\beta+1}$, such that $c_\alpha \in V_{\theta_{\alpha+1}}$ and $\theta_\alpha \in c_{\alpha+1}$ and $c_{\alpha+1}$ forces another specific ordinal into $\dot{D}$ above θ_α. By Σ_2-elementarity, there is such an element below $\theta_{\alpha+1}$. It follows that the final limit condition $c = c_{\omega\beta+1}$ forces that $\dot{D}$ is unbounded in every $\theta_{\omega\alpha}$ for $\alpha \leq \beta$ and consequently that $\{\, \theta_{\omega\alpha} \mid \alpha \leq \beta \,\}$ is a closed $\beta{+}1$ sequence in the intersection $\dot{D} \cap \check{T}$. Thus, c forces this instance of fat stationarity, and so T remains fat stationary in $\mathcal{V}[C]$, as desired. $\square$

The class club-shooting forcing will be central to the proof of the main theorem. In that context, our fat-stationary class S will itself be Cohen generic, a case for which one can mount an easier direct proof of distributivity. Meanwhile, we thought it worthwhile to provide the general analysis here.

6. The class Fodor principle is not provable in KM

The main theorem follows from the results proved in this section.

Theorem 13. *Every countable model of Kelley-Morse set theory $\mathcal{V} = \langle V, \in, \mathscr{C} \rangle \models \mathrm{KM}$ has an extension $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle \models \mathrm{KM}$, adding new classes but no new sets, with a regressive class function $F : \mathrm{Ord} \to \omega$ that is not constant on any stationary class.*

Proof. Assume that $\mathcal{V} = \langle V, \in, \mathscr{C} \rangle$ is a countable model of KM. We shall perform forcing in several steps and then take a symmetric submodel of the final model. For the first step of forcing, let $F : \mathrm{Ord} \to \omega$ be a $\mathcal{V}$-generic Cohen function, added over this model with the class forcing notion $\mathbb{P}$ consisting of set initial segments of F. Since this forcing is κ-closed for every cardinal κ, it follows that the forcing is tame and adds no new sets. Consider the corresponding forcing extension $\mathcal{V}[F] = \langle V, \in, \mathscr{C}[F] \rangle \models \mathrm{KM}$.

For each $n < \omega$, let $A_n = \{\, \gamma \in \mathrm{Ord} \mid F(\gamma) = n \,\}$ be the pre-image of n. We claim that this is a fat-stationary class in the extension $\mathcal{V}[F]$. To see this, consider any ordinal β and suppose that $\dot{C}$ names a class club. Below any condition $f_0 \in \mathbb{P}$, we may strengthen it to force another specific element into $\dot{C}$ above the domain of f_0, and by iterating this ω-many times, we may find a condition f_1 extending f_0 and forcing unboundedly many elements into $\dot{C} \cap \mathrm{dom}(f_1)$. If $\gamma = \mathrm{dom}(f_1)$, we are now free to extend further to a condition f_1^+ for which $f_1^+(\gamma) = n$. This condition therefore forces that γ is in $A_n \cap \dot{C}$. By now iterating this process β many times, we find a condition f_β^+ that forces a closed copy of $\beta + 1$ into $A_n \cap \dot{C}$. So A_n is fat stationary in $\mathcal{V}[F]$.

Let $B_n = \{\, \gamma \in \mathrm{Ord} \mid F(\gamma) > n \,\}$, which is the same as $\bigcup_{m>n} A_m$, the union of the tails of the A_m's. Thus, each B_n is the union of fat-stationary classes and consequently fat stationary itself. In $\mathcal{V}[F]$, let $\mathbb{Q}_n$ be the class-club-shooting forcing for B_n. Now let $\mathbb{Q} = \prod_{n<\omega} \mathbb{Q}_n$ be the full-support product in $\mathcal{V}[F]$ of all these club-shooting forcing notions.[1]

For any $n < \omega$, let $\mathbb{Q} \restriction n = \prod_{m \leq n} \mathbb{Q}_m$ be the finite product of the $\mathbb{Q}_m$. This forcing can be seen as first shooting a club into the fat-stationary class B_0, and then into B_1, and so on, up to B_n. By lemma 12, each of these classes remains fat-stationary in the next extension, since $B_0 \supseteq B_1 \supseteq \cdots \supseteq B_n$, and so by induction the product $\mathbb{Q} \restriction n$ is $<$Ord-distributive, tame, and adds no sets. (We could alternatively argue directly that $\mathbb{P} * \dot{\mathbb{Q}} \restriction n$ has a dense subclass that is κ-closed for every κ.)

Suppose that $G \subseteq \mathbb{Q}$ is $\mathcal{V}[F]$-generic for the full product forcing, and consider the resulting extension $\mathcal{V}[F][G]$. This will not be a model of KM, nor even of GBC,

---

[1] Everything in the following argument would have worked just as well had we used finite-support instead of full-support.

since it has the sequence of clubs $C_n \subseteq B_n$, added by $\mathbb{Q}$, as a coded sequence, but the intersection $\bigcap_n C_n$ must be empty, since any γ in that intersection has no possible value for $F(\gamma)$.

Our desired final model, however, will not be $\mathcal{V}[F][G]$, but rather a certain symmetric submodel of this model, which we shall argue is a model of KM. Namely, let $\mathscr{C}^+$ consist of the classes added by some stage of the $\mathbb{Q}$ forcing, that is, the classes in $\mathcal{V}[F][G \upharpoonright n]$, for some $n < \omega$. Our desired final model is $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle$.

We claim that $\mathcal{V}^+$ is a model of KM. Notice first that each of the finite-product models $\mathcal{V}[F][G \upharpoonright n]$ is a model of KM with the same sets as V, because $\mathbb{P} * \mathbb{Q} \upharpoonright n$, we argued above, is $<$Ord-distributive. It follows that $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle$, being the union of a chain of KM models, all with the same first-order part, is a model at least of GBC. This is because any finitely many classes of $\mathscr{C}^+$ exist together in some finite extension $\mathcal{V}[F][G \upharpoonright n]$, where we get all the desired instances of first-order class comprehension using those class parameters.

It remains to see that $\mathcal{V}^+$ is fully a model of KM. For this, we need to prove the second-order class comprehension axiom. Suppose that $Z \in \mathscr{C}^+$ and consider the class $\{\, x \mid \mathcal{V}^+ \models \varphi(x, Z) \,\}$, where φ may have second-order class quantifiers. The parameter Z exists in some $\mathcal{V}[F][G \upharpoonright n]$. Since this is a model of KM, the forcing relation for the rest of the product forcing $\mathbb{Q}^n = \prod_{m \geq n} \mathbb{Q}_m$, augmented with the unary predicate $\check{\mathscr{C}}$ as discussed in section 4, is definable in $\mathcal{V}[F][G \upharpoonright n]$. Using the predicate $\check{\mathscr{C}}$ we can express in the augmented forcing language that a class Z is in $\mathscr{C}^+$ by saying that Z is the interpretation of an element of $\check{\mathscr{C}}$ by $G \upharpoonright n$ for some n. Furthermore, the product forcing $\mathbb{Q}^n$ is locally homogeneous in $\mathcal{V}[F][G \upharpoonright n]$ by local isomorphisms that act separately on each coordinate using the homogeneity isomorphism from the proof of lemma 12; the resulting local isomorphism of $\mathbb{Q}$ therefore respects each $\mathcal{V}[F][G \upharpoonright n]$ and consequently does not affect $\mathscr{C}^+$. Because of this, the question of whether $\varphi(x, Z)$ holds in $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle$ is settled by all conditions in the same way. It follows that the class $\{\, x \mid \mathcal{V}^+ \models \varphi(x, Z) \,\}$ exists already as a class in $\mathcal{V}[F][G \upharpoonright n]$, definable there using the forcing relation, and consequently this class is in $\mathscr{C}^+$. So $\mathcal{V}^+$ is a model of KM.

Finally, notice that in this model, each B_n contains a class club, and furthermore, the sequence of B_n exists as a coded class $B = \{\, (n, \gamma) \mid \gamma \in B_n \,\}$, since this class is definable directly from F. Meanwhile, $\bigcap_{n<\omega} B_n$ is empty, since for any ordinal γ we have $\gamma \notin B_n$ when $n = F(\gamma)$. $\square$

Corollary 14. *Kelley-Morse set theory* KM*, if consistent, does not prove that the club filter is σ-closed. It is relatively consistent with* KM *that there is a coded sequence $\langle B_n \mid n < \omega \rangle$ of classes, each in the club filter, but the intersection $\bigcap_n B_n = \emptyset$ is empty.*

That situation is exactly what the model in theorem 13 provides.

The main theorem of this article subsumes a previous result of the first two authors. Specifically, in unpublished work [GH] using similar methods as here, Gitman and Hamkins had constructed symmetric models of Kelley-Morse set theory KM in which the class choice principle failed even in very weak formulations, establishing the following as their main result, a result which can now be seen as a consequence of the main result of this article.

Corollary 15 (Gitman, Hamkins [GH])**.** *Kelley-Morse* KM *set theory, if consistent, proves neither the class choice principle* CC*, nor the principle* $\mathrm{CC}(\Pi^0_1)$ *nor even the principle* $\mathrm{CC}_\omega(\Pi^0_1)$*.*

Proof. Theorem 7 shows that these class choice principles each imply the Fodor principle for regressive functions $F : \mathrm{Ord} \to \omega$, but theorem 13 shows that this can fail with KM. □

Corollary 16. *The theory* KM *plus the negation of the Fodor principle, and hence also the negation of* CC *and even* $\mathrm{CC}_\omega(\Pi^0_1)$*, is conservative over* KM *for first-order assertions about sets.*

Proof. Theorem 13 shows that every countable model of KM has an extension, with the same first-order part, in which the Fodor principle fails. So any first-order statement failing in a model of KM also fails in a model of KM plus the negation of the class Fodor principle. □

This method therefore provides a way of *softly* killing KM + CC, in the killing-them-softly sense of [Car15]. In particular, since tame forcing preserves KM + CC by theorem 11, this shows that we should expect some kind of symmetric extension to be needed, rather than a full forcing extension, in order to kill the class choice principle CC.

7. Separating the class Fodor principles

In this section, we shall separate the various Fodor principles we defined in the introduction.

Theorem 17. *Every countable model* $\mathcal{V} = \langle V, \in, \mathscr{C}\rangle \models \mathrm{KM} + \mathrm{CC}$ *has an extension to a model* $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+\rangle \models \mathrm{KM}$*, with new classes but no new sets, in which the class* $<$Ord*-Fodor principle holds and indeed* CC_κ *holds for every cardinal* κ*, but the class Fodor principle and the bounded class Fodor principle fail.*

Proof. Assume that $\mathcal{V} = \langle V, \in, \mathscr{C}\rangle$ is a model of KM+CC. We shall follow the proof of theorem 13, but here we begin by adding a generic regressive function $F : \mathrm{Ord} \to \mathrm{Ord}$, rather than $F : \mathrm{Ord} \to \omega$. Let $\mathbb{P}$ be the class forcing to add such a function F via bounded conditions, and consider the forcing extension $\mathcal{V}[F] = \langle V, \in, \mathscr{C}[F]\rangle$. Since this forcing is tame, this is a model of KM + CC by theorem 11.

For any ordinal ξ, let $B_\xi = \{\, \gamma \mid F(\gamma) > \xi \,\}$, in analogy with the classes B_n in the proof of theorem 13, and define $\mathbb{Q}_\xi$ to be the class-club-shooting forcing for B_ξ. Consider the set-support product $\Pi_{\xi\in\mathrm{Ord}}\mathbb{Q}_\xi$, and let $\mathbb{Q} \restriction \alpha = \Pi_{\xi\leq\alpha}\mathbb{Q}_\xi$ be initial segments of $\mathbb{Q}$.

We argue that $\mathbb{P} * \mathbb{Q} \restriction \alpha$ has a dense subclass that is κ-closed for every cardinal κ. Specifically, let $\mathbb{D}$ be the subclass of $\mathbb{P} * \dot{\mathbb{Q}} \restriction \alpha$ consisting of conditions that are pairs $(p, \check{q})$, such that all the closed sets appearing in q have the same supremum γ, which is also the maximal element in the domain of p, and such that $p(\gamma) > \alpha$. This is easily seen to be dense, since any condition $(p, \dot{q})$ can be extended to a condition with a check name in the second coordinate, since the $\mathbb{P}$ forcing does not add sets, and then one can extend all the closed bounded sets appearing in q to have the same supremum and arrange such a γ and $p(\gamma)$ as desired. This class is κ-closed for every κ, simply by taking unions in each coordinate and defining the value of p at the new maximal element to be above α. Thus, $\mathbb{P} * \mathbb{Q} \restriction \alpha$ is $<$Ord-distributive and consequently tame, and it adds no new sets.

Let $G \subseteq \mathbb{Q}$ be $\mathcal{V}[F]$-generic and consider the submodel of $\mathcal{V}[F][G]$ of the form $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle$, where $\mathscr{C}^+$ are the classes added by some stage of the forcing. Analogous arguments to those given in the proof of theorem 13 show that $\mathcal{V}^+$ satisfies KM and that the bounded class Fodor principle fails there, since the function F is not bounded on any stationary proper class. Consequently, the class Fodor principle also fails.

It remains to show that CC_κ holds in $\mathcal{V}^+$ for every cardinal κ, which by theorem 7 implies <Ord-Fodor. Suppose that $\mathcal{V}^+ \models \forall\alpha{<}\kappa\, \exists X\, \varphi(\alpha, X, Z)$ with a class parameter Z. Since $Z \in \mathcal{V}^+$, there must be a stage γ such that $Z \in \mathcal{V}[F][Q \restriction \gamma]$, which is a model of KM + CC by theorem 11. Consequently, $\mathcal{V}[F][G \restriction \gamma]$ has a forcing relation for the tail forcing $\mathbb{Q}^\gamma$ in the language with the ground-model class predicate $\check{\mathscr{C}}$, used to define $\mathcal{V}^+$ inside the full extension. So there is some condition $p \in G$ forcing that for every $\alpha < \check{\kappa}$, there is a class X in $\mathcal{V}^+$ such that $\varphi^{\mathcal{V}^+}(\alpha, X, \check{Z})$. For each $\alpha < \kappa$ there is an open dense subclass of conditions forcing that a specific name $\dot{X}$ for some specific extension $\mathbb{Q} \restriction \beta$ witnesses $\varphi(\check{\alpha}, \dot{X}, \check{Z})$. Since the forcing is <Ord-distributive, we may intersect these open dense classes and find a single condition p^+ extending p and forcing specific class names as witnesses. By the class-choice principle CC_κ in $\mathcal{V}[F][G \restriction \gamma]$, we may choose for each $\alpha < \kappa$ a specific such class name $\dot{X}_\alpha$ and ordinal β_α, such that $p^+ \Vdash \varphi^{\mathcal{V}^+}(\check{\alpha}, \dot{X}_\alpha, Z)$, with the κ-sequence of such names and ordinals coded in $\mathcal{V}[F][G \restriction \gamma]$. From this κ-sequence of names $\dot{X}_\alpha$, it is easy to construct the name of the corresponding κ-sequence, which is added by stage $\sup_{\alpha<\kappa} \beta_\alpha$, and so p^+ forces this instance of CC_κ in $\mathcal{V}^+$, as desired. □

Essentially the same arguments give the following.

Theorem 18. *Suppose $\mathcal{V} = \langle V, \in, \mathscr{C} \rangle$ is a countable model of* KM + CC *and κ is a regular cardinal in V. Then $\mathcal{V}$ has an extension $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle \models \mathrm{KM} + \mathrm{CC}_\lambda$, for every $\lambda < \kappa$, adding new classes but no new sets. In this model the class κ-Fodor principle fails, but the class λ-Fodor principle holds for every $\lambda < \kappa$.*

Theorem 19. *Every countable model $\mathcal{V} = \langle V, \in, \mathscr{C} \rangle \models \mathrm{KM} + \mathrm{CC}$ has an extension to a model $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle \models \mathrm{KM}$, adding new classes but no new sets, in which the bounded class Fodor principle holds, but the class ω-Fodor principle fails.*

Proof. We use the model $\mathcal{V}^+ = \langle V, \in, \mathscr{C}^+ \rangle$ from the proof of theorem 13, and also use all the notation from that proof. The difference is that we now also assume CC in $\mathcal{V}$. It suffices to show that the bounded class Fodor principle holds in $\mathcal{V}^+$. Let $H : S \to \mathrm{Ord}$ be a regressive function on a stationary class S in $\mathcal{V}^+$. The class H is added by some stage $\mathcal{V}[F][G \restriction m]$ for some $m < \omega$. Let $S_\alpha = \{\, \gamma \mid H(\gamma) < \alpha \,\}$, and note that $\alpha < \beta \to S_\alpha \subseteq S_\beta$. If no S_α is stationary in $\mathcal{V}^+$, then for each α there is stage $n = n_\alpha$ at which S_α becomes non-stationary in $\mathcal{V}[F][G \restriction n]$. By homogeneity, this n_α does not depend on G and the function $\alpha \mapsto n_\alpha$ exists already at stage m. So there must be some n occurring for unboundedly many α. So unboundedly many S_α and hence all S_α are nonstationary already at this stage $\mathcal{V}[F][G \restriction n]$. But this is a model of CC, and consequently satisfies the class Fodor theorem, contradicting that H is not bounded on any stationary class in this model. □

8. The class Fodor principle is invariant by set forcing

In this section, we shall show that the class Fodor principle is preserved by set forcing and that set forcing cannot make the class Fodor principle hold in a forcing extension if it did not already hold in the ground model. The main preliminary fact is that after set-sized forcing, every new class club contains a ground model club. This is because if $p \Vdash \dot{C}$ is a class club, then the class of ordinals α such that $p \Vdash \check{\alpha} \in \dot{C}$ is closed and unbounded, by the class analogue of the usual argument for κ-c.c. forcing and club subsets of κ.

Theorem 20. *The class Fodor principle is invariant by set forcing over models of* GBC^-. *That is, it holds in an extension if and only if it holds in the ground model.*

Proof. Assume GBC^-, and consider a set-forcing extension $\mathcal{V}[G] = \langle V[G], \in, \mathscr{C}[G]\rangle$ of $\mathcal{V} = \langle V, \in, \mathscr{C}\rangle$, where $G \subseteq \mathbb{P} \in V$ is $\mathcal{V}$-generic.

Suppose first that the class Fodor principle holds in the ground model $\mathcal{V}$, and suppose that $F : S \to \mathrm{Ord}$ is a class function in the extension $\mathcal{V}[G]$ that is regressive on the stationary class $S \subseteq \mathrm{Ord}$ in $\mathscr{C}[G]$. Let $\dot{S}$ and $\dot{F}$ be class $\mathbb{P}$-names for which $S = \dot{S}_G$ and $F = \dot{F}_G$. In the ground model $\mathcal{V}$, for each ordinal α choose a condition $p_\alpha \in \mathbb{P}$ forcing $\check{\alpha} \in \dot{S}$, if possible, and deciding the value of $\dot{F}(\check{\alpha})$. Since S is stationary in the extension and therefore meets all ground-model class clubs there, it follows that p_α is defined on a stationary class of ordinals. Since $\mathbb{P}$ is a set, it follows by the class Fodor principle that $p_\alpha = p$ is constant on a stationary class S_0 of ordinals. For $\alpha \in S_0$, let $E(\alpha)$ be the value of $\dot{F}(\check{\alpha})$ that is forced by p. This is a regressive function on S_0 in the ground model, and so by the class Fodor principle, it is constant on a stationary subclass $S_1 \subseteq S_0$. So p forces that $\dot{F}$ is constant on $\check{S}_1$, which remains a stationary class in $\mathcal{V}[G]$, since as we observed before the theorem, every new class club contains a ground-model class club. So p forces this instance of the class Fodor principle in $\mathcal{V}[G]$. In particular, by relativizing below any given condition, there can be no condition forcing a violation of the class Fodor principle in $\mathcal{V}[G]$, and so it must hold there, as desired.

Conversely, assume that the class Fodor principle holds in a set-forcing extension $\mathcal{V}[G] = \langle V[G], \in, \mathscr{C}[G]\rangle$ of $\mathcal{V} = \langle V, \in, \mathscr{C}\rangle$, and that $F : S \to \mathrm{Ord}$ is a regressive class function on a stationary class S in the ground model. Since every new class club contains a ground model class club, S remains stationary in $V[G]$, and so by the class Fodor principle in $\mathcal{V}[G]$, the function F is constant on a stationary class in $\mathcal{V}[G]$. So $F^{-1}(\{\alpha\})$ is a stationary class in $\mathcal{V}[G]$ for some ordinal α. But this class exists already in the ground model, and therefore must be stationary there. So the class Fodor principle holds in $\mathcal{V}$, as desired. □

The downward absoluteness of the class Fodor principle can be easily generalized to Ord-c.c. class forcing, for which also every new class club contains a ground model class club.

Question 21. *Is the class Fodor principle invariant by pre-tame class forcing? Is it preserved by such forcing? By* Ord*-c.c. class forcing?*

9. Fodor on higher class well orders

Let us consider an analogue of Fodor on class well-orders beyond Ord. Namely, if Γ is a class well-order of uncountable cofinality, then we have sensible notions of *club in* Γ and *stationary in* Γ, as well as notions of *regressive* functions on Γ.

First, let's notice that class well orders have cofinality at most Ord.

Lemma 22. *Assume* GBC*. Every class well-order* Γ *contains a closed unbounded suborder of type* Ord *or type* κ *for some regular cardinal* κ*.*

Proof. Suppose that Γ is a class well order. If some set is unbounded in the Γ order, then we can find a cofinal κ-sequence for some regular cardinal κ. So assume that no set is unbounded in the Γ order. In this case, let a_0 be the least element of Γ, and define $a_{\alpha+1}$ to be the Γ-least element above a_α and above all elements of V_α, and a_λ to be the supremum of a_α for $\alpha < \lambda$, when λ is a limit ordinal. It follows that the class $\{\, a_\alpha \mid \alpha \in \mathrm{Ord}\,\}$ is an unbounded subclass of Γ of order type Ord. $\square$

Thus, every class well-order Γ has a *cofinality* $\kappa \leq \mathrm{Ord}$, which is the smallest order-type of any closed unbounded subclass.

Let us briefly note that in GB, even without the global axiom of choice (or AC for sets), one may always take the underlying class of a proper class well-order Γ to be Ord. The reason is that if $\Gamma = \langle X, \lhd \rangle$ is a proper class well order, then the underlying class X is stratified by $\in$-rank. So we can enumerate the class X in a set-like well-ordered manner, enumerating first by rank and then by the Γ-order within each rank. This provides a bijection of X with Ord, and so Γ is isomorphic to a relation on Ord.

Theorem 23. *Assume* GBC *and the class Fodor principle. Then for any class well-order* Γ*, every regressive class function* $F : S \to \Gamma$ *defined on a stationary subclass* $S \subseteq \Gamma$ *is bounded on a stationary subclass* $T \subseteq S$*.*

Proof. Suppose that $F : S \to \Gamma$ is regressive on a stationary class $S \subseteq \Gamma$. By lemma 22, there is a closed unbounded subclass $C \subseteq \Gamma$ of order type $\kappa \leq \mathrm{Ord}$. It is easy to see that $S \cap C$ remains stationary in Γ. Define a class function $F_0 : C \to C$ $F_0(a) = \sup\{\, b \in C \mid b \leq F(a)\,\}$, which presses the value of $F(a)$ down to C. The class function F_0 is regressive on C, and is therefore isomorphic to a regressive function on κ. By the class Fodor principle (needed only for the case $\kappa = \mathrm{Ord}$), it follows that $F_0 \restriction T$ is constant for some stationary subclass $T \subseteq S \cap C$. The class T remains stationary in Γ, and it follows that $F(a)$ is pressed down always to the same element of C for $a \in T$. Thus, $F(a)$ is bounded by the next element of C for $a \in T$, and so F is bounded on a stationary subclass of S. $\square$

One cannot expect in general that regressive functions on class well-orders exceeding Ord will be constant on a stationary subclass, since if Γ is a class well-order exceeding Ord, then by lemma 22, there is a closed unbounded subclass $C \subseteq \Gamma$ of order-type at most Ord, and we can define $F(\gamma)$ to be the Γ-order type of the Γ-predecessors of γ in C, that is, the set $\{\, \alpha \in C \mid \alpha <_\Gamma \gamma\,\}$. This order-type is strictly less than Ord, and so the function F is regressive on the part of Γ beyond Ord, but it cannot be constant on any unbounded subclass of Γ.

10. Questions on the strength of the class Fodor principle

In addition to question 21, we have many further questions about the class Fodor principle. To our knowledge all the questions asked in this article are completely open.

First, although the main theorem shows that KM does not prove the class Fodor principle, nevertheless we do not know much about how the Fodor principle is situated in the hierarchy of second-order set theories.

Question 24. *What is the strength of the class Fodor principle in the hierarchy of second-order set theories?*

For example, if we suspect that the class Fodor principle is strong, then we inquire:

Question 25. *Does* GBC *plus the class Fodor principle imply the principle* ETR*? Does it imply the existence of a first-order truth predicate? Does it imply* Con(ZFC)*?*

The folklore argument at the end of the introduction shows that in any ω-standard model of GBC with the class Fodor principle, there must be a stationary class of κ for which $V_\kappa \models \mathrm{ZFC}$. This is evidence in favor of strength for the class Fodor principle, but it doesn't quite seem to answer the question.

If we suspect, in contrast, that the class Fodor principle is weak, then we inquire:

Question 26. *Is* GBC *plus the class Fodor principle conservative over* ZFC *for first-order assertions in set theory?*

It is not true that every GBC model extends to a model of the class Fodor principle, since the model of our main theorem is a model of GBC with a class function $F : \mathrm{Ord} \to \omega$, which is not constant on any stationary class, and this would continue to be a violation of the class Fodor principle in any GBC extension. This can be seen as evidence against an affirmative answer to question 26.

Question 27. *Suppose* $0^\sharp$ *exists, and consider the* GBC *model* $\langle L, \in, \mathrm{Def}(L)\rangle$*, where* $\mathrm{Def}(L)$ *has only the first-order definable classes with set parameters. Does this model satisfy the class Fodor principle?*

This model has some weak forms of Fodor. Namely, if $F : S \to \mathrm{Ord}$ is definable from the indiscernible parameters $i_{\xi_0}, \ldots, i_{\xi_n}$ and some indiscernible i_ξ with $\xi > \xi_n$ is in S, then F is constant on a stationary subclass. In this case, all indiscernibles above i_ξ must also be in S, and moreover, F has the same value on all of them. But there are stationary classes in this model, such as Cof^L_ω, which do not contain any indiscernibles, so that observation does not settle the question.

The question is related to stationary reflection. Namely, if $S \cap \kappa$ is stationary in κ for some large enough κ, where $L_\kappa \prec L$, then consider $F \restriction \kappa$. There must be $\alpha < \kappa$ with $F^{-1}(\{\alpha\})$ actually stationary in κ, hence definably stationary in κ. Apply elementarity. So F has constant value α on a stationary class.

References

[AFG] Carolin Antos, Sy David Friedman, and Victoria Gitman. Boolean-valued class forcing. *manuscript under review.*

[Ant18] Carolin Antos. Class forcing in class theory. In *The hyperuniverse project and maximality*, pages 1–16. Birkhäuser/Springer, 2018.

[Car15] Erin Carmody. *Force to change large cardinal strength.* ProQuest LLC, Ann Arbor, MI, 2015. Thesis (Ph.D.)–City University of New York.

[Fri00] Sy D. Friedman. *Fine structure and class forcing*, volume 3 of *de Gruyter Series in Logic and its Applications.* Walter de Gruyter & Co., Berlin, 2000.

[GH] Victoria Gitman and Joel David Hamkins. Kelley-Morse set theory and choice principles for classes. In preparation.

[GH16] Victoria Gitman and Joel David Hamkins. Open determinacy for class games. In Andrés E. Caicedo, James Cummings, Peter Koellner, and Paul Larson, editors, *Foundations of Mathematics, Logic at Harvard, Essays in Honor of Hugh Woodin's 60th Birthday*, AMS Contemporary Mathematics. 2016. Newton Institute preprint ni15064.

[GHH+17] Victoria Gitman, Joel David Hamkins, Peter Holy, Philipp Schlicht, and Kameryn Williams. The exact strength of the class forcing theorem. *ArXiv e-prints*, July 2017. manuscript under review.

[GHJ16] Victoria Gitman, Joel David Hamkins, and Thomas A. Johnstone. What is the theory ZFC without Powerset? *Math. Logic Q.*, 62(4–5):391–406, 2016.

[HKS18] Peter Holy, Regula Krapf, and Philipp Schlicht. Characterizations of pretameness and the Ord-cc. *Ann. Pure Appl. Logic*, 169(8):775–802, 2018.

[Kar18] A. Karagila. Fodor's lemma can fail everywhere. *Acta Math. Hungar.*, 154(1):231–242, 2018.

[MM75] W. Marek and A. Mostowski. On extendability of models of ZF set theory to the models of Kelley-Morse theory of classes. pages 460–542. Lecture Notes in Math., Vol. 499, 1975.

[Neu51] Walter Neumer. Verallgemeinerung eines satzes von alexandroff und urysohn. *Mathematische Zeitschrift*, 54(3):254–261, Sep 1951.

[Wil18] Kameryn J. Williams. *The Structure of Models of Second-Order Set Theories*. ProQuest LLC, Ann Arbor, MI, 2018. Thesis (Ph.D.)–City University of New York.

(V. Gitman) The City University of New York, CUNY Graduate Center, Mathematics Program, 365 Fifth Avenue, New York, NY 10016

Email address: vgitman@nylogic.org

URL: https://victoriagitman.github.io/

(J. D. Hamkins) Professor of Logic, Oxford University & Sir Peter Strawson Fellow, University College, Oxford. High Street Oxford OX1 4BH U.K.

Email address: joeldavid.hamkins@philosophy.ox.ac.uk

URL: http://jdh.hamkins.org

(A. Karagila) University of East Anglia. Norwich, NR4 7TJ, UK

Email address: karagila@math.huji.ac.il

URL: http://karagila.org/