φ-Features at the Syntax-Semantics Interface: Evidence from Nominal Inflection

Ivona Kučerová

I argue for a novel model of feature valuation in the CI interface and explore under what circumstances a syntactic feature is semantically interpretable. As the groundwork for the investigation, I propose an explicit Distributed Morphology model of Italian nouns of profession. The data provide evidence that the morphology accesses the narrow-syntax representation at two different temporal points within a phase: the earlier point (Spell-Out) returns a morphological realization faithful to feature values present in narrow syntax, while the later point (Transfer) allows for a narrow-syntax representation to be enriched by the CI component. Thus, there is no syntactic distinction between interpretable and uninterpretable features: a syntactic feature appears to be interpretable only if it has been licensed by the CI interface.

Keywords: syntax-semantics (CI) interface, φ-features, Agree, agreement

While it has been recognized in the syntax and semantics literature that gender and number features within a DP may be valued either from the syntactic structure (typically from the lexicon) or from the context (e.g., Corbett 1991, Dahl 2000, Wechsler and Zlatić 2000, 2003, Sauerland 2004, Wiltschko and Steriopolo 2007, Acquaviva 2008, Heim 2008, Wiltschko 2009, Spathas 2010, Sudo 2012, Matushansky 2013, Merchant 2014), little is known about how these two types of valuation are established or what structural properties underpin this distinction. Furthermore, there is a growing consensus that more than one structural element can be the source of the valued features and that the structurally higher element tends to be responsible for the semantically informed value (e.g., √ and n in Kramer 2009; see Pesetsky 2013 and Landau 2016 for related configurational analyses). Yet it is far from clear what it means for features to be valued from the "context" or why the structural height of functional heads should make a difference with respect to whether or not a syntactic feature will be semantically interpreted.

This article brings new data from Standard Italian into the debate and uses them to argue for a novel treatment of semantically interpretable φ-features, specifically gender. The Standard Italian nominal system morphologically marks two distinct genders, three distinct nominal classes (idiosyncratic nominal endings), and two numbers, a combination that lends itself to a theoretical investigation of complex gender interactions. Crucially, Italian nominal inflection also distinguishes between grammatical (i.e., idiosyncratic) and natural (i.e., context-dependent) gender.

The main empirical focus of the article is on names of professions that exhibit morphosyntactically varied behavior: for some speakers, the root *chirurg-* 'surgeon' can denote a female even if it is grammatically masculine (*il chirurgo* 'the.M surgeon'[1]), while other speakers prefer a derived form with a distinct morphological feminine marker, be it a feminine determiner or a combination of a feminine determiner and a feminine ending on the noun (*la chirurgo* 'the.F surgeon' vs. *la chirurg-a* 'the.F surgeon-F'). Some nouns, such as *avvocato* 'lawyer', display the same range of formal variation, even though there is a designated feminine derived noun (*avvocatessa* 'female lawyer'). The structural variability in turn allows us to investigate the timing of related syntactic operations and their interactions in a controlled way.

The article is organized as follows. In section 1, I present an explicit Distributed Morphology model of Italian nominal morphosyntax. I provide evidence that Italian nouns come in three distinct flavors, one group being lexically specified for gender in the lexicon, and the other two being unspecified. I show that type of gender specification has morphosyntactic consequences. This section provides the necessary groundwork for the main theoretical proposal presented in section 2. There, I elaborate on recent insights of phase theory (Chomsky 2001, 2008, 2013, Narita 2011) and insights from the formal semantic literature on the interpretation of gender features on pronouns (e.g., Heim 2008, Sudo 2012). I argue for a novel model of feature valuation at the syntax-semantics interface and explore under what circumstances a syntactic feature may be semantically interpretable. I show that the morphology interface can access the narrow-syntax representation at two different points of the derivation within a phase: the earlier point—technically, Spell-Out—returns a morphological realization that is faithful to feature values present in narrow syntax, while the later point—technically, Transfer—allows a narrow-syntax representation to be enriched by the conceptual-intentional (CI) component. Thus, there is no syntactic distinction between interpretable and uninterpretable features: a syntactic feature appears to be interpretable only if it has been licensed by the CI interface.

1 How Many Genders Are There?

In his foundational study of Spanish nominal inflection, Harris (1991) provides extensive evidence that gender and class markers must be modeled as two separate entities.[2] The same pattern that motivated Harris's conclusion arises in Standard Italian (e.g., Dressler and Thornton 1996, Thornton 2001, 2003a,b, Ferrari 2005, Ferrari-Bridgers 2007, 2008). In general, native speakers of Italian share the intuition that vocalic endings on nominals correspond to grammatical gender:[3] namely, a noun that ends in *-o* in the singular is grammatically masculine (M) (where *grammatically masculine* means that it triggers masculine agreement on agreeing elements such as determiners or predicative adjectives), while nouns ending in *-a* are grammatically feminine (F), and nouns

---

[1] Throughout the article, I mark gender in glosses on elements that agree in gender with the noun in question. Abbreviations used in the article: M: masculine, F: feminine, GEN: gender, NUM: number, SG: singular, PL: plural, CL: class.

[2] I use the term *class marker* as a descriptive label for idiosyncratic nominal vowel endings. The theoretical status of these markers is one of the questions to be addressed in this study.

[3] Unless indicated otherwise, the data are from a project with Anna Moro. The data were collected either from native speakers or from systematic dictionary, grammar book, and Internet searches.

ending in -*e* can be either. Even though this correlation is frequent, a closer look at the data does not support this intuition beyond its being a statistical tendency.[4] While some nouns ending in -*o* are indeed grammatically always masculine, such as *il libro* 'the.M book', a handful of -*o* nouns are grammatically always feminine, such as *la mano* 'the.F hand'. Some -*o* nouns denoting people may be grammatically masculine or feminine depending on the natural gender of their referent, such as *il soprano*, *la soprano* 'the.M male soprano/sopranist, the.F soprano' and *il chirurgo*, *la chirurgo* 'the.M surgeon, the.F surgeon'.[5]

The same type of gender variation is attested with nouns ending in -*a* and -*e*: for instance, *l'artista bravo* 'the-artist good.M' vs. *l'artista brava* 'the-artist good.F', *il giornalista* 'the.M journalist' vs. *la giornalista* 'the.F journalist', and *il preside* 'the.M head (of something)' vs. *la preside* 'the.F head (of something)', *il vigile* 'the.M sentry' vs. *la vigile* 'the.F sentry'.

From these data, we can conclude that class and grammatical gender cannot correspond to the same grammatical representation because the same ending can correspond to either gender and both genders can be realized by any of the vocalic endings. It follows that if both class and gender are part of the morphosyntactic module,[6] they must be modeled as separate entities.

Interestingly, there is a class of nouns that challenges this conclusion, namely, Harris's (1991) so-called *mating nouns*. Here, the vocalic ending is strictly determined by the *natural* gender of the referent. More precisely, if a mating noun ends in -*o*—for instance, *il bambino*, *il ragazzo*—then it denotes a male or an underspecified individual. Thus, *bambino* denotes a baby boy or a baby, while *ragazzo* denotes a boy or a child in general. In contrast, -*a* versions of mating nouns strictly denote females: thus, *la bambina* 'the baby girl', *la ragazza* 'the girl'.[7]

The attested pattern thus raises the following questions: (a) What is the structural relationship between grammatical and natural gender? (b) What is the relationship between class marker and gender? (c) Why does natural gender "rewrite" the nominal ending in some instances, while leaving it intact in others?

1.1 Dissociating Class Marker and Gender

We have seen that we need to dissociate the grammatical realization of gender and class marker (or at least the nominal vocalic endings) in order to account for the basic data patterns in Standard

[4] According to Ferrari-Bridgers (2007), 32% of Italian nouns in her database are feminine -*a* nouns and 32% are masculine -*o* nouns. Authors differ in how they treat "exceptions." While for Thornton (2001, 2003a,b) the exceptions are significant and require a separate treatment of class and gender, Ferrari (2005; Ferrari-Bridgers 2007, 2008) believes that because the gender of 98% of Italian nouns can be predicted by hierarchically ordered rules (phonological, morphological, and semantic), gender does not need to be represented in the structure. Others, most notably Lampitelli (2010, 2014), assume that there is an isomorphic mapping between gender and vowel endings; that is, they do not acknowledge the existence of exceptions. As we will see, however, not distinguishing between gender and class misses an important empirical generalization.

[5] The agreement properties of nouns that can be grammatically feminine or masculine depending on the natural gender of their referent are complicated by various sociolinguistic factors, such as social prestige. I will discuss the variation in ending and agreement properties in detail in section 1.3.

[6] I assume a Distributed Morphology style of grammar architecture (Halle and Marantz 1993); however, I will use the term *morphosyntactic* as a useful shorthand for syntactic features and representations that play a role in the morphological realization module.

[7] See Percus 2011 for an extensive discussion of such examples.

Italian.[8] I argue that class is a strictly morphophonological reflex of the root lexical specification (analogically to the treatment of theme vowels in Oltra-Massuet 1999, Embick and Halle 2005, Embick and Noyer 2007), while the gender feature (as well as number) corresponds to a formal syntactic feature. Later we will need to distinguish between gender that is an idiosyncratic part of the lexical specification and gender that is introduced at the syntax-semantics interface (i.e., valued from the context), but for now the discussion concerns only grammatical gender.

If a class value and/or gender value is idiosyncratic, then it needs to be associated with the root already in the lexicon, be it a valued feature of some sort or a lexical diacritic.[9] Formally, either the root is stored in the lexicon as a nominal structure (i.e., a structure already including category-defining n: $[_n \; n_{[CL/GEN]} \; \sqrt{root}]$), or the root comes with a valued class feature/diacritic that will restrict which functional head will be able to combine with it (i.e., $\sqrt{root}_{[CL/GEN]}$). As for the latter option, either n would be merged in syntax (as category-defining n) and the class and/or gender feature value would be assigned via Agree between the root and n, or a realization rule would be needed that would make insertion of the vocalic ending sensitive to the properties of the root.[10] The morphological realization of the noun then arises via some form of morphological affixation (be it m-merger, affix hopping, or whatever bound morpheme procedure the reader prefers) and morphological mapping based on the Subset Principle (Halle and Marantz 1993). The mapping rules for n need to reflect number, which we do not have in the current toy system; the mapping rules in (1) will suffice for expository purposes.

(1) a. $[\text{CL}:1, \text{GEN}:\alpha] \rightarrow$ -o
      b. $[\text{CL}:2, \text{GEN}:\alpha] \rightarrow$ -a
      c. $[\text{CL}:3, \text{GEN}:\alpha] \rightarrow$ -e

Since the system in place treats gender and class features as two separate variables, any gender feature can combine with any class feature, which is indeed the case. Thus, we straightforwardly derive the possibility of having masculine -*o* nouns, masculine -*a* nouns, feminine -*o* nouns, feminine -*a* nouns, masculine -*e* nouns, and feminine -*e* nouns.

To obtain a complete DP, merger of D and Num is required (Ritter 1993, 1995, Borer 2005). Since D agrees in number and gender, it will need to have unvalued NUM and GEN features. For concreteness, let us assume that n comes from the lexicon with a valued number feature as well—a vast simplification, but it will do for present purposes as number is not central to the investigation.[11]

---

[8] The system to be developed here departs from most of the existing literature for empirical and theoretical reasons. To my knowledge, none of the existing proposals takes into account the full range of data discussed here. Furthermore, none of the work investigates either the derivational consequences of the proposed structures to be discussed in section 1.4 or the syntactic-semantic consequences to be discussed in section 2.

[9] In Embick 2000, Embick and Noyer 2007, and subsequent work, it is argued that even though the phonological form of a root is not syntactically represented, other properties, such as features, might in principle be accessible to syntax. See also Lowenstamm 2008 for an extensive argument isolating gender as being either on $\sqrt{}$ or very close to it.

[10] This would mean, however, that technically such a root would be category-neutral. In section 1.4, we will see evidence that these roots are indeed category-specific.
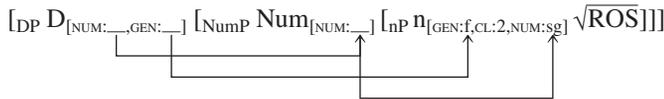
[11] See, for instance, Alexiadou 2011 and Ouwayda 2014 for a sophisticated argument that number, analogically to gender, may be valued in more than one place in the structure.

I use simplified labels F/M and SG/PL in order to avoid for now the question of the exact feature geometry.

I assume downward-probing Agree, where matching and valuation are distinct operations (Chomsky 2000, Adger 2003, Pesetsky and Torrego 2007).[12] The complete step-by-step derivation of *la rosa* 'the.F rose' is given in (2).

(2) la rosa *'the.F rose': feminine -a noun*

   a. Feature distribution taken from the lexicon and feature matching

$$[_{DP} \, D_{[NUM:\_\_,GEN:\_\_]} \, [_{NumP} \, Num_{[NUM:\_\_]} \, [_{nP} \, n_{[GEN:f,CL:2,NUM:sg]} \, \sqrt{ROS}]]]$$

    b. Feature valuation by Agree and morphological output

$$[_{DP} \, [_{D[NUM:sg,GEN:f]} \, \text{la}] \, [_{NumP} \, Num_{[NUM:sg]} \, [_{nP} \, [_{n[GEN:f,CL:2,NUM:sg]} \, \text{-a}] \, \text{ros}]]]$$

The same structure gives us a masculine *-o* noun such as *il libro* 'the.M book'—that is, another case in which the vocalic ending seems to correspond to gender—simply by changing the value of the gender feature into M and the class into 1. In order to derive a feminine *-o* noun such as *la mano* 'the.F hand', we need to alter only the value of gender on n. A simplified derivation is given in (3). The boxes indicate which features are central to the morphological realization.[13]

(3) la mano *'the.F hand': feminine -o noun*

$$[_{DP} \, [_{D} \, \text{la}_{[NUM:sg,\boxed{GEN:f}]}] \, [_{nP} \, \text{-o}_{[GEN:f,\boxed{CL:1},NUM:sg]} \, \text{man}]]$$
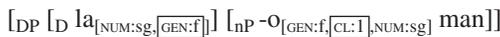
The proposed model strictly separates gender and class. Even though they both need to be represented in the structure, as they are crucial for morphological insertion, they fundamentally differ in that only gender interacts with the rest of the syntactic structure, simply because higher functional heads do not contain an unvalued class feature. Consequently, class does not interact with the structure beyond nP.[14]

1.2 Natural Gender and the Order of Syntactic Operations

We have already seen data suggesting that the gender of some nouns is not determined from the lexicon, as it may depend on the natural gender of their referent. This group is exemplified by

---

[12] With the system set up this way, gender and number features need to probe separately because if we included a requirement on ϕ-completeness, nothing would ever get valued within a DP. (See Danon 2011 for a detailed discussion of the issue of ϕ-completeness within DP.) Since the exact representation of number is not the focus of this article, I will remain agnostic about whether this is an appropriate representation. Note also that the proposed model crucially assumes that nominal concord relations, at least the relation between D and its complement, are based on Agree. But see Norris 2014 for an alternative view.

[13] For clarity of presentation, the boxes do not include the number features even though the number value is crucial for the morphological realization as well.

[14] The system thus differs crucially from other proposals for Standard Italian—most notably, those of Ferrari (2005; Ferrari-Bridgers 2007, 2008) and Lampitelli (2010, 2014)—that assume that class is some PF feature but do not make a strict dissociation between class and gender. Even though these systems do well with the subset of data they concentrate on, they do not extend to the more complex patterns to be discussed in this article.

-o nouns that can be grammatically either masculine or feminine, such as *il/la soprano* 'the.M/ F soprano' and *il/la chirurgo* 'the.M/F surgeon'. The same gender duality accompanied by no change in the form of the noun itself is attested with *-a* nouns (e.g., *artista* 'artist' and *autista* 'driver') and *-e* nouns (e.g., *preside* 'head (of something)' and *vigile* 'sentry'). The vocalic ending of these nouns is strictly determined by their class, but their gender may be determined by their natural referent.[15]

Furthermore, we have seen that some nouns lack a class specification as well, as not only the gender of their determiner but also their vocalic ending is determined by the natural gender of their referent. These are the so-called mating nouns such as *il bambino* 'the.M baby/baby boy' vs. *la bambina* 'the.F baby girl', *il ragazzo* 'the.M kid/boy' vs. *la ragazza* 'the.F girl'. Table (4) summarizes the possible combinations.

(4)	Type	Class?	Gender from lexicon?
	il libr-o/la ros-a	✓	✓
	il/la chirurg-o	✓	*
	il bambin-o/la bambin-a	*	*

These empirical facts suggest that roots may come with a varied degree of nominal specification: (a) some roots are valued for class and gender, (b) some roots are valued only for class, and (c) some roots have no nominal specification (i.e., these roots are category-neutral). I will call these three basic types of roots (a) conservative, (b) semiconservative, and (c) liberal.[16] The corresponding structures are given in (5).[17]

(5) *Possible root representations taken from the lexicon*
　　a. Conservative (*il libr-o/la ros-a*): [$_n$ n$_{[GEN,CL]}$ √root]
　　b. Semiconservative (*il/la chirurg-o*): [$_n$ n$_{[CL]}$ √root]
　　c. Liberal (*il bambin-o/la bambin-a*): √root

Furthermore, the distributional facts suggest that if gender is not valued from the lexicon, it is valued from the context. Putting aside for now what it means to be "valued from the context" and how such an operation could be technically implemented, it is useful to see what consequences this intuitive parameterization has for the proposed syntactic structure. If we assume that the lexically valued gender feature is part of narrow syntax while the contextually determined value becomes available only at the syntax-semantics interface, we predict that when an Agree operation is valued by the lexically specified gender feature, this valuation must happen *before* the gender feature can be valued by its contextual counterpart.

---

[15] There is an asymmetry in the semantic interpretation of the gender on the determiner: while the grammatically masculine form (*il chirurgo*, *il soprano*) may denote a male or a female (*il chirurgo* being more often used to denote a male referent, *il soprano* a female), their feminine counterparts must denote a female. We will see in section 2.3 how the gender markedness follows from the theory of failed Agree combined with implicated presuppositions.

[16] The naming convention reflects the fact that historically, all Italian nouns had a theme vowel. The emergence of classless nouns is a relatively new innovation.

[17] Theoretically, there could also be a gender-only root specification: [$_n$ n$_{[GEN]}$ √root]. As we will see, such a representation is not empirically attested.

Consequently, we predict the order of syntactic operations to be as in (6). For concreteness, I assume that the locus of the contextual valuation is D; since D is the phase head (or at least the labeling head in the extended lexical domain), it is the only functional head locally accessible to the CI interface at the point of Spell-Out. The proposed ordering follows from the mechanics of Agree. No additional stipulation is needed.
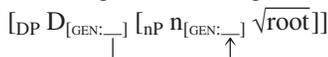
(6) *The timing of syntactic operations*
  1. Merge of D
  2. Probing of φ-features on n by φ-features on D
  3. Matching
  4. Valuation by Agree using features already present in the structure
  5. Valuation of D from the context for the features that have not yet been valued
  6. Automatic valuation by Agree of remaining features within the previously established matching link
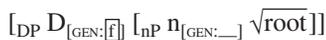
Since the gender feature on D and the gender feature on n are part of the same Agree chain, D and n cannot undergo gender feature valuation independently of each other. Valuation via D is possible only if the gender feature on n has not been valued. Thus, depending on whether or not the root comes with a valued gender feature, we obtain two basic gender-feature-valuation configurations within a DP.

(7) *Configuration I: Gender from the context (D)* (il bambino, la bambina; il/la chirurgo)
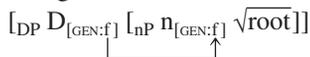 a. Matching of unvalued gender features
    $[_{DP} D_{[GEN:\_\_]} [_{nP} n_{[GEN:\_\_]} \sqrt{root}]]$
  b. D cannot be valued by n because the gender feature on n is not valued → D valued from the context
    $[_{DP} D_{[GEN:\boxed{f}]} [_{nP} n_{[GEN:\_\_]} \sqrt{root}]]$
  c. The gender feature on n is automatically valued via the existing matching link with D
    $[_{DP} D_{[GEN:f]} [_{nP} n_{[GEN:f]} \sqrt{root}]]$

(8) *Configuration II: Gender from the lexicon (n)* (il libro, la rosa)
  a. Matching of gender features on D and n
 $[_{DP} D_{[GEN:\_\_]} [_{nP} n_{[GEN:\boxed{f}]} \sqrt{root}]]$
  b. Valuation of the gender feature on D by Agree with the valued gender feature on n
 $[_{DP} D_{[GEN:f]} [_{nP} n_{[GEN:f]} \sqrt{root}]]$
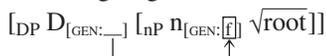
Each of these configurations could in principle have a class feature on the root as well. Let us first consider configuration I in (7), that is, the structure with no gender on the root. If the root has a class feature, we get a noun of the *il/la chirurgo* 'the.M/F surgeon' type. If the root has no nominal specification, we get a mating noun, such as *il bambino/la bambina* 'the.M baby/

baby boy'/'the.F baby girl'. Simplified derivations of the former type—*il/la chirurgo*—are given in (9) and (10). The only difference between the derivations is that in (10), the gender feature on D is valued as feminine. In this configuration, agreement with the rest of the syntactic structure is strictly determined by the value of the gender feature introduced by D.

(9) il chirurgo *'the.M surgeon'*
   a. Feature distribution from the lexicon and matching
      $[_{DP} D_{[GEN:\_\_]} [_n n_{[GEN:\_\_,CL:1]} \sqrt{CHIRURG}]]$
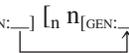
   b. Feature valuation and morphological output
      $[_{DP} [_D{}_{[\boxed{GEN:m}]} il] [_n n_{[GEN:m,\boxed{CL:1}]} $ -o] chirurg]]
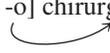
(10) la chirurgo *'the.F surgeon'*
   $[_{DP} [_D la_{[\boxed{GEN:f}]}] [_{nP} -o_{[GEN:f,\boxed{CL:1}]} $ chirurg]]

Before we can see how mating nouns are derived, we need to expand the set of morphological mapping rules for n. The reason is that the vocalic ending on these nouns is strictly determined by their contextually determined gender but the mapping rules strictly require a class value. Thus, we need to add morphological mapping rules for cases in which there is no class feature. The expanded set of morphological mapping rules is given in (11). As before, these rules leave out number.

(11) *Mapping rules for n*
   a. $[CL:1, GEN:\alpha] \leftrightarrow$ -o
   b. $[CL:2, GEN:\alpha] \leftrightarrow$ -a
   c. $[CL:3, GEN:\alpha] \leftrightarrow$ -e
   d. $[GEN:m] \leftrightarrow$ -o
   e. $[GEN:f] \leftrightarrow$ -a

With such mapping rules and the current system in place, we straightforwardly derive the alternation attested in mating nouns.

(12) a. il bambino *'the.M baby/baby boy'*
      $[_{DP} [_D il_{[\boxed{GEN:m}]}] [_{nP} -o_{[\boxed{GEN:m}]} bambin]]$
   b. la bambina *'the.F baby girl'*
      $[_{DP} [_D la_{[\boxed{GEN:f}]}] [_{nP} -a_{[\boxed{GEN:f}]} bambin]]$

1.3 Valuation from the Context vs. Valuation from the Lexicon

As mentioned in footnote 5, names of professions display a persistent gender and agreement variation in Italian. In most cases, the variation exists because sociological prestige has traditionally been associated with maleness even if the profession now includes women. Good examples are *l'avvocato* 'the.M lawyer', *il chirurgo* 'the.M surgeon', and *il ministro* 'the.M minister'. We see a parallel grammatical variation with *il soprano* 'the.M soprano', which—although grammatically

masculine—typically denotes a female singer but may denote a male singer, especially a boy, as well.[18]

Feminine versions of masculine nouns in Standard Italian may be formed with the feminine derivational suffix *-essa*, as in *avvocatessa* 'female lawyer'. Some speakers strongly prefer this formation. For other speakers, however, the feminine counterparts can be derived by inflectional morphology as well. That is, the original grammatical gender value from the lexicon is reanalyzed by some speakers so the noun's denotation can accommodate current female uses. If we concentrate on the inflectional formation only, there are three attested forms that match the root specifications given in (5), repeated here.

(13) *Possible root representations taken from the lexicon*
    a. Conservative: [$_n$ n$_{[GEN:m,CL:1]}$ √root]
    b. Semiconservative: [$_n$ n$_{[CL:1]}$ √root]
    c. Liberal: √root

Speakers whose lexicon contains option (13c) thus produce *il ministro, la ministra* 'the.M minister, the.F minister', *il soprano, la soprana* 'the.M male soprano/sopranist, the.F soprano', *l'avvocato, l'avvocata* 'the.M lawyer, the.F lawyer'. Consequently, since D is valued only for natural gender, any agreeing elements in the structure will agree with this gender: for example, *il soprano bravo* 'the.M male soprano/sopranist good.M' and *la soprana brava* 'the.F soprano good.F'.

Speakers who prefer the semiconservative option, (13b), produce *il soprano* and *la soprano*—that is, nouns that do not store gender in the lexicon yet retain their class marker. Again, as with mating nouns, for such nouns the gender value on D is determined by their referent, and agreeing syntactic elements agree with this gender: for example, *il soprano bravo* 'the.M male soprano/sopranist good.M' and *la soprano brava* 'the.F soprano good.F'.[19]

Finally, there are speakers who treat nouns like *il chirurgo* as grammatically masculine; that is, all elements within the extended DP are morphologically marked as masculine. These speakers seem to prefer the conservative option, (13a). Surprisingly, this seemingly grammatically masculine DP triggers feminine agreement on predicates, a fact the system developed here so far cannot account for. Thus, while prescriptive grammars require *avvocato* to be grammatically masculine even if it refers to a female, feminine agreement is common, as these Internet examples attest: *brava avvocato . . . una donna meravigliosa* 'good.F lawyer.M . . . one.F woman wonderful.F' (Facebook; accessed 16 April 2014); *brava avvocato, grazie di causa* 'good.F lawyer.M, thanks to (the) case' (Twitter, February 2013). The same fact is acknowledged by other sources, such as the *Enciclopedia Italiana di scienze, lettere ed arti*, which gives the following examples: *il soprano è andato* 'the.M male soprano/sopranist is gone.M', *il soprano è andata* 'the.M male

---

[18] Historically, *il soprano* is subject to the same sociological shift, as female voices in operas were typically performed by male singers until the end of the eighteenth century.

[19] As far as I have been able to establish, all adjectives have the same ϕ-features as the determiner; that is, we do not find a distinction between low and high adjectives of the sort reported by Pesetsky (2013) for gender in Russian and by Landau (2016) for number in Hebrew.

soprano/sopranist is gone.F', and *la soprano è andata* 'the.F soprano is gone.F'. The data show that both the determiner and the participle can be masculine (*il . . . andato*), or they both can be feminine (*la . . . andata*), or the determiner can be masculine and the participle feminine (*il . . . andata*).

The attested interspeaker variation thus confirms the proposed model with three distinct representations of nominal roots. However, the mixed agreement pattern attested with the conservative option suggests that grammatical and natural gender can cooccur within the same DP. In section 2, I will develop a model of feature valuation that will capture the full range of facts.

1.4 Predictions: Category-Neutral Roots

Before we investigate feature valuation, it is useful to consider three robust predictions the current system makes to ensure that the proposed implementation is, in principle, on track. First, it predicts that some roots should come from the lexicon being specified for nominal features, while others should seem to be category-neutral. Since only mating nouns are based on category-neutral roots (as they do not carry any idiosyncratic marking from the lexicon), only mating nouns should have corresponding verbs derived from the same root with no special verbal derivational morphology—that is, derived only by the default inflectional suffix *-are*. This prediction is borne out. As far as I have been able to establish, there are no verbs based on the roots of nouns that carry a class marker, as approximated in (14).[20] In contrast, mating nouns tend to have verbal counterparts, as shown in (15), based on examples from Fabrizio 2013.[21]

(14) *Noun-verb pairs based on category-specific roots are nonexistent*
   a. il libro 'the book' ↔ *librare 'to book'
   b. la rosa 'the rose' ↔ *rosare 'to rose'

(15) *Noun-verb pairs based on category-neutral roots are attested*
   a. la sposa 'the bride', lo sposo 'the groom' ↔ sposare 'to marry'
   b. il figlio 'the son', la figlia 'the daughter' ↔ figliare 'to give birth'
   c. l'astrologo, l'astrologa 'astrologer' ↔ astrologare 'to practice astrology'
   d. il monaco 'the monk', la monaca 'the nun' ↔ monacare 'to put someone into a convent, to make him/her become a monk/nun'
   e. il commissario, la commissaria 'commissioner' ↔ commissariare 'to put under a commissioner'

Two more predictions concern loanwords. The class of a loanword and its related inflectional properties are often derived from its phonotactic shape (see, e.g., Pesetsky and Torrego 2011 for a discussion of undeclinable loanwords in Russian). If a loanword ends in a consonant, it cannot

---

[20] Some of these pairs might be excluded on semantic grounds. What matters, though, is that I am not aware of *any* verb based on the same root as a class-marked noun (putting aside denominal verbs derived by specialized verbal morphology).

[21] Not all mating nouns have a verbal counterpart. This is presumably because of additional restrictions, such as the semantic interpretation of the root. What matters is that there are some that do.

be associated with any class marker, as association with a class marker is strictly based on a word-final vowel. Such a loanword cannot be associated with gender from the lexicon either, as Italian does not have nouns that come from the lexicon with grammatical gender but no class. Consequently, consonant-final loanwords should be based on category-neutral roots. Furthermore, if they denote objects, they cannot have their gender assigned by D, as objects do not have a contextually determined gender. As a result, they should be realized with default grammatical gender, namely, masculine. The system thus predicts that consonant-final loanwords should be masculine and that they should be able to have verbal counterparts derived without specialized derivational morphology. Both of these predictions are borne out: for example, *chat* 'the.M chat' and *chattare* 'to chat', *il film* 'the.M film' and *filmare* 'to film', *il blog* 'the.M blog' and *bloggare* 'to blog'.

1.5 Summary

To summarize, we have seen empirical evidence that Italian nouns come in three different structural configurations: they can come from the lexicon with class and grammatical gender, with class but no grammatical gender, or with no class and no gender. If a noun comes from the lexicon with no grammatical gender, its gender can be determined from the context. The nature of context valuation is explored in the next section.

2 Features at the Syntax-Semantics Interface

The data reviewed in section 1 belong to a larger family of phenomena that display a difference between so-called grammatical and semantic/natural gender (e.g., Corbett 1991, Sauerland 2004, Wiltschko and Steriopolo 2007, Neeleman 2008, Matushansky 2013), with some authors noting an additional difference between lexical and referential gender (Dahl 2000, Wechsler and Zlatić 2000, 2003). A parallel distinction has been made in the semantics literature regarding the ф-feature specification of pronouns and other anaphoric elements. The relevant observations are that pronominal agreement and gender specification on pronouns is anaphoric, but the source of the relevant antecedent is in principle twofold (Heim and Kratzer 1998): the ф-feature valuation can be determined either from the structure (syntax) or from the context; and while some gender features are semantically interpreted, others are only morphological reflexes of some sort of feature transmission (e.g., von Stechow 2003, Heim 2008, Kratzer 2009, Sudo 2012). In my analysis, I follow the essence of the semantic approaches and tie their findings to the syntactic analysis proposed in section 1.

The core hypothesis I test here is that either the value of the gender feature on D can be determined by Agree with a gender feature valued from the lexicon, or it can be determined from the context. If the gender feature on D is valued from the context, it is always semantically interpreted because the value restricts the natural gender of the referent, as illustrated by Dahl's (2000:106) example, based on Corbett's (1991:181) discussion of common gender: *The doctor said he/she could see me tomorrow.* If the valued feature comes from the lexicon, the empirical picture is complicated by the fact that the lexically determined gender often coincides with the

denotation implied by the lexical semantics of the root. That is, it is not clear a priori whether *la donna* 'the.F woman' denotes a female referent by virtue of the lexical semantics of its root, the feminine gender feature determined from the lexicon, or both. I will adopt the stronger position and argue that the gender feature valued from the lexicon is never semantically interpreted.[22] If there is a natural-gender meaning associated with a DP, the semantic interpretation always comes from elsewhere.

Let us start by making the following assumption: for the morphological form of a definite article to depend on a contextually determined gender, the gender must in some way be represented in the structure of the DP; otherwise, the morphological component would not know whether to insert *il* or *la*, and so on. Once we make this assumption, three questions arise: First, where is this formal representation located? Second, if it is a formal feature, how is it valued? Third, if the morphology module reflects a contextual/semantic valuation, is the valuation compatible with the Y-model of grammar?

I explore the first and second questions in sections 2.1 and 2.2, respectively. As for the third, this grammar architecture consideration is not unique to gender. It also arises in scope and morphological marking phenomena that are sensitive to semantic and pragmatic considerations (number, definite vs. indefinite articles, quantifier raising, scrambling, etc.). The implementation proposed here belongs to the family of works that argue for an economy condition that prefers isomorphic relations between interface representations. Specifically, if the two interface representations differ, then we observe a derivational asymmetry that favors the LF representation over its PF counterpart (e.g., Williams 2003, Reinhart 2006, Bobaljik and Wurmbrand 2012, Kučerová 2012). That is to say, the morphological realization is sensitive to the semantically informed output representation. In section 2.3, I explore how the primacy of the semantic representation arises.

2.1 Phase Transfer as the Source of Contextually Determined Valuation

We have already seen, especially in section 1.4, that at least for some Standard Italian nouns gender is determined on the root through the lexicon.[23] However, if the gender value is contextually dependent, it cannot be determined on the root, unless we posit two distinct lexical entries—one for feminine, one for masculine.[24] If the contextually determined gender is not on the root, it needs to be located somewhere within the extended nominal projection, either as a valued feature or as a dedicated functional projection consisting of the feature. This type of dissociation has been argued for by Sauerland (2004), Wiltschko and Steriopolo (2007), Kramer (2009, 2015), Wiltschko (2009), Steriopolo and Wiltschko (2010), Matushansky (2013), and Pesetsky (2013),

---

[22] In syntactic terms, this amounts to saying that there is no distinction between interpretable and uninterpretable features, at least not for gender, as the gender feature is always uninterpretable.

[23] By *on the root*, I mean within the nominally specific structure stored in the lexicon, in the sense discussed in section 1.
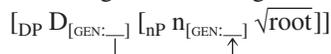
[24] An analysis that posits two separate lexical entries predicts that any gaps in distribution would be accidental. As we saw in section 1, the distribution is structurally restricted.

among others. The authors differ with regard to where the other gender feature is located, the two main candidates being n (most prominently in Kramer 2009, 2015) and D (as in Sauerland 2004 and Steriopolo and Wiltschko 2010).[25] As for the context-sensitive gender value, I argue that, at least in Standard Italian, it is assigned on D.
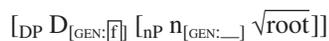
The syntactic structures for the two distinct places of valuation proposed in section 1.2, (7)–(8), are repeated here.

(16) *Configuration I: Gender from the context (D)* (il bambino, la bambina)
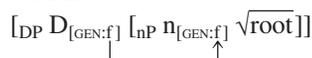   a. Matching of unvalued gender features
      [$_{DP}$ D$_{[GEN:\_\_]}$ [$_{nP}$ n$_{[GEN:\_\_]}$ $\sqrt{root}$]]
   b. D cannot be valued by n because the gender feature on n is not valued → D valued from the context
      [$_{DP}$ D$_{[GEN:\boxed{f}]}$ [$_{nP}$ n$_{[GEN:\_\_]}$ $\sqrt{root}$]]
   c. The gender feature on n is automatically valued via the existing matching link with D
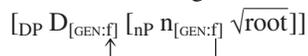   [$_{DP}$ D$_{[GEN:f]}$ [$_{nP}$ n$_{[GEN:f]}$ $\sqrt{root}$]]

(17) *Configuration II: Gender from the lexicon (n)* (il libro, la rosa)
   a. Matching of gender features on D and n
      [$_{DP}$ D$_{[GEN:\_\_]}$ [$_{nP}$ n$_{[GEN:\boxed{f}]}$ $\sqrt{root}$]]
   b. Valuation of the gender feature on D by Agree with the valued gender feature on n
      [$_{DP}$ D$_{[GEN:f]}$ [$_{nP}$ n$_{[GEN:f]}$ $\sqrt{root}$]]
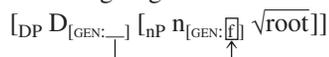
The core assumption is that a formal syntactic feature can be assigned a context-sensitive value only at the syntax-semantics interface. More precisely, the only point during the syntactic derivation when a formal syntactic feature may be valued by the CI component is when the phase is sent to the syntax-semantics interface (so-called Transfer). For concreteness, I model valuation by the CI component as minimal search in the sense of Chomsky (2008, 2013) and Narita (2011) and argue that contextual valuation takes place as part of labeling of the transferred phase. Since the label is determined by features of the phase head, only unvalued features on D may be valued by the CI interface at Transfer. Crucially, if such a phase is selected by a higher syntactic structure, the feature valued by the CI interface will now be visible for Agree within the narrow-syntax derivation and may in turn value other instances of formal syntactic features.

If the locus of semantic feature valuation is a phase head, one might wonder whether such a feature could be assigned to an n head instead, assuming that nP might be a phase as well (Kramer 2009, 2015). The idea is that only a complete DP can interact with natural gender because

---

[25] A notable exception is Wiltschko 2009 and subsequent work, where it is argued that languages differ with regard to the exact location of the contextually determined feature, and whether or not the corresponding functional head is syntactically an adjunct.

the assignment of natural gender is tied to establishing a referential address (via PERSON,[26] as we will see in section 2.2). In other words, the semantic properties of DP provide a formal trigger for minimal search from the CI interface.[27]

In contrast, nP does not have a semantic representation that requires an association with a referential address. Properties have only grammatical gender.[28] In turn, the semantic component cannot associate nP with a natural gender feature. Formally, there is no feature on n that would trigger minimal search by the CI interface. While nP might be a complete unit from the point of view of narrow-syntactic derivation, only DP forms a semantically complete unit, that is, a structure that can be sent to the CI module (Chomsky 2008, 2013, Larson 2011, Arsenijević and Hinzen 2012).

In the next section, I argue that contextual valuation is dependent on PERSON, a feature that is visible both to the narrow syntax and to the CI interface. Before I proceed to the investigation of gender and PERSON, however, two notes are in order. First, one might wonder whether the proposed system utilizes some notion of semantic agreement. Though it might superficially seem to do so, so far it is based strictly on the syntactic notion of Agree.[29] Second, the proposal is intellectually indebted to the insights of Wechsler and Zlatić (2000, 2003), who to my knowledge were the first to make a clear distinction among class ("declension"), valuation within narrow syntax ("concord"), valuation from the context ("index"), and semantic interpretation ("semantics"). Some distinctions between their proposal and mine are attributable solely to differences in architecture between Head-Driven Phrase Structure Grammar and the Minimalist Program. Others are more fundamental, however, such as differences in the lexicon representation and the formal dependency between gender and PERSON to be explored next.

2.2 Licensing PERSON at the CI Interface

We currently have two sources of gender valuation: the gender value of a DP can be either determined idiosyncratically from the lexicon or supplied by D. Crucially, D seems to provide a gender value only if it refers to an animate referent—more precisely, to a human or anthropomor-

---

[26] I use small capitals to systematically denote PERSON as a syntactic feature. I do not adopt the same convention for gender and number as numerous instances of these terms in the article simultaneously refer to the syntactic object and its morphological realization.

[27] The relevant structure typically corresponds to an individual (semantic type e). However, as an anonymous reviewer points out, restricting the semantic requirement to individuals and in turn to referential interpretation is too strong, as it predicts that indefinites and quantifiers never receive natural gender interpretation. This is not the case, as we find reflexes of natural gender with indefinites and some quantifiers. For instance, among professions that do not really exist, one participant in an online discussion of pervasive gender inequality in contemporary Italian society lists *nessuna chirurga* 'no.F surgeon.F' (a post from 6 September 2008, http://www.mentecritica.net/matteo-73/oldstuff/comandante-nebbia /6496/). Although I do not discuss QPs and indefinites here, I use the notions of index (modeled as a variable and receiving its value via an assignment function in the sense of Heim and Kratzer 1998) and referential address (in the sense of referential-like properties of quantifiers as discussed in Endriss 2009) to avoid this incorrect prediction.

[28] The reason is that they do not have a D layer with a PERSON feature that would associate with an index.

[29] The proposed system might seem reminiscent of Collins and Postal's (2012) idea that there might be a pronoun embedded within a syntactically larger DP. Here, however, φ-features do not have two distinct values within the same DP. The purpose of modeling D as a bundle of φ-features is to formalize how natural gender values might be made accessible to syntax proper.

phous referent.[30] To understand this correlation, we first need to understand the properties of gender with respect to other syntactic features. I argue that contextually determined gender depends on properties of another syntactic feature, namely, PERSON (e.g., Harley and Ritter 2002, Béjar and Rezac 2003, Adger and Harbour 2007). Furthermore, I argue that the syntactic dependency has a semantic counterpart: the interpretive difference between a grammatical and a contextually determined gender feature results from the *association* of gender and PERSON, instead of the gender feature per se.

The central role of PERSON for the syntax-semantics interface has been argued for in both the syntax and the semantics literature, starting with the foundational work of Ritter (1995), who to my knowledge was the first to associate PERSON with the D head. More recently, Longobardi (2008), Landau (2010), and Sudo (2012), among others, have argued for a connection between PERSON and a semantic (referential) index. I follow this line of work and argue that PERSON is the core narrow-syntax feature that plays a role in mapping ɸ-features onto the semantic representation. More precisely, if we model it as [±participant] (Nevins 2007 and the literature cited there), PERSON is the most obvious candidate for a syntactic counterpart of a semantic index (i). Technically, an index is a numerical index (i.e., a variable that is mapped to semantic values by an assignment function; Heim and Kratzer 1998).[31] The crucial assumption here is that although a ɸ-feature in narrow syntax might have a semantic counterpart, the semantic object does not have to be identical to the syntactic object. Informally, we can understand the connection as an argument-tracking device that ties the narrow-syntax representation to its semantic counterpart. The question then is how the formal connection between PERSON and the index representation arises, and how the gender feature interacts with it.

Before we turn to the syntactic representation of the association, let us unwrap the interpretive part. Sauerland (2004) and Matushansky (2013) independently propose that valuation of context-dependent gender features is driven by the semantic component. The crucial insight is that semantic treatments of gender do not interpret gender directly; instead, they treat it as a *presupposition* associated with an assignment index (in simplified terms, as a pointer to the actual referent, or a referential address). Consider Dahl's (2000) example again.

(18) The doctor said he/she could see me tomorrow.

The linguistic antecedent *doctor* is compatible with both a female and a male referent, yet the pronoun used to refer back to the antecedent must respect the natural gender of the referent. The idea is that the form of the pronoun is determined by the index. The gender value of the pronoun restricts the interpretation of the antecedent: if the pronoun is masculine, the referent is masculine; if the pronoun is feminine, the referent is feminine. This intuitive correlation can be formally

---

[30] I use the term *anthropomorphous* somewhat loosely to have a cover term for instances of animate objects that grammatically display natural-gender properties.

[31] Using variable representations avoids some nontrivial difficulties associated with constructing context-dependent valuation as Agree. Technically, as we will see shortly, binding ɸ-features in syntactic terms corresponds to two separate processes: valuation of PERSON with respect to an index, and association of the index with gender and person presuppositions that may require overt morphological realization in order to satisfy Heim's (1991) Maximize Presupposition principle.

captured as an admissibility condition on the referent, that is, as a presupposition associated with the gender features. The corresponding semantic denotation of masculine and feminine gender is given in (19), modeled after Heim 2008. The denotation of the features is defined as an identity function of type $\langle e, e \rangle$. That is, the function takes an individual as its argument and returns an individual only if the individual is of the appropriate gender. If the individual is not of the appropriate gender, the function remains undefined and the structure is not interpretable.[32]

(19) a. $[\![[\text{GEN}{:}\text{f}_i]]\!]^{w,g} = \lambda x_e . \ g(i)$ is female in w: x
       b. $[\![[\text{GEN}{:}\text{m}_i]]\!]^{w,g} = \lambda x_e . \ g(i)$ is a person in w: x

The semantic formula asserts that a feminine feature associated with index i denotes a female person only if the referent associated with the index is indeed a female. For its masculine counterpart, the formula asserts that such an index denotes a person only if the individual the index points to is indeed a person.

Since the interpretation function in (19) ultimately interprets an assignment index (i) associated with the gender, not the actual gender feature, it will return a semantic denotation corresponding to natural gender only if PERSON is associated with a human or anthropomorphous referent in the CI module. By definition, a gender feature is not semantically interpreted in and of itself.[33] Its interpretive effect lies solely in imposing interpretation restrictions on an index associated with PERSON. Thus, the dependency of humanness we observed in the Italian examples is a side effect of the licensing of PERSON by the CI interface.

While (19a) specifies feminine grammatical gender as denoting a female, masculine gender is in principle compatible with both natural genders. This is exactly what we find in Standard Italian. Thus, feminine *la chirurgo* 'the.F surgeon' must denote a female, whereas *il chirurgo* 'the.M surgeon' can denote a person of either gender. That the masculine form of such a noun is often interpreted as referring to a male, even though the masculine feature presupposes a person instead of a male person, results from an implicated presupposition in the sense of Sauerland (2003, 2008). Consequently, if the implicated presupposition of maleness conflicts with the lexical semantics of the root (as, e.g., in *il soprano* 'the.M soprano'), the implicated presupposition does not arise and the grammatically masculine form freely refers to a female. Finally, since only gender features that presuppose a (female) person give rise to a natural-gender interpretation, the proposed semantic system explains why the natural-gender interpretation arises only for human and anthropomorphous referents.[34]

With a better understanding of the interpretive properties of the gender feature, we can now turn to the syntactic representation. Two major questions arise: (a) How is the association between PERSON and the semantic index created? and (b) How can the semantic component value a syntactic gender feature within a DP?

---

[32] The subscript i in the denotation of a gender feature indicates that the admissibility restriction arises only in the context of a semantic index.

[33] Though it might obtain a natural gender interpretation indirectly from the lexical semantics of the root.

[34] If we instead assumed that some instances of the syntactic gender feature are interpretable and some are uninterpretable (as, e.g., in Smith 2015), the connection to humanness would need to be stipulated. Here, the connection follows from independently needed assumptions.

Regarding the first question, I follow Chomsky's (2013, 2015) insight and argue that it happens as part of the labeling of a DP phase. According to Chomsky, Merge is a set-forming operation that in and of itself does not provide the newly created syntactic object with a label. Instead, labeling is a procedure triggered by the semantic interface.[35] In an endocentric structure, the closest head becomes the label.[36] In the nominal cases investigated here, the label of a DP corresponds to the D head—more precisely, to the bundle of features that forms the D head. I argue that since the labeling of such a DP is triggered by the semantic interface, it is in the labeling process that a semantic index is added to the label as well.

We can now turn to the second question: how can the association of PERSON with an index by the CI interface value a gender feature within a DP? That is, how can the information provided by the CI module be relevant for syntax proper and determine the value of the syntactic feature (gender)? As has been extensively argued in the Distributed Morphology literature, morphological realization may, but does not have to, faithfully realize features present in the syntactic representation. Feature bundles may be *adjusted* at the interface, and Vocabulary Insertion may be sensitive to the structural environment. Under this view, morphological gender marking can arise as the representation of PERSON in a particular structural context. I argue that the context-dependent valuation of the gender features in the structure arises from such an adjustment. Specifically, I argue that the context-dependent valuation of gender features falls under a large family of morphosyntactic phenomena governed by Heim's (1991) Maximize Presupposition principle, which asserts that if there is a presupposition associated with a structure and if this presupposition can be grammatically realized, it must be realized.[37] Since semantic gender features are presuppositional, gender-feature-sensitive morphological insertion as a *last resort* is governed by Maximize Presupposition as well.

To see how Maximize Presupposition–driven morphological realization works, let us look at English pronouns. Consider again example (18), repeated here.

(20) The doctor said he/she could see me tomorrow.

The analysis of pronouns proposed by Postal (1969) and Elbourne (2005) argues that pronouns are morphological realizations of DPs with an elided NP. If we model D as a bundle of unvalued φ-features, the morphological realization of the pronoun reflects the φ-features on D valued by Agree with valued φ-features of the unpronounced NP. The problem is that unless we posit that English nouns are valued for grammatical gender and that there are two separate lexical entries for nouns like *doctor*—one grammatically feminine, one grammatically masculine—there is no valued gender feature on the NP. In turn, even if D came with an unvalued gender feature from

---

[35] See Narita 2011 and Chomsky 2013, 2015 for arguments, primarily from criterial freezing structures, that labeling is required by the semantic module.

[36] Technically, the head is identified by so-called minimal search. The actual label is formed by an identity merge of the features present in the head.

[37] Maximize Presupposition has been argued to account for a variety of semantically sensitive morphosyntactic phenomena: for instance, the distribution of definite articles (Heim 1991), number marking (Sauerland 2003, 2008), tense marking (Sauerland 2002), and scrambling (Kučerová 2012).

the lexicon—which is far from obvious—there is no valued gender feature on the goal. For concreteness, I follow Ritter 1995 and much subsequent work in that I model PERSON as a valued feature directly introduced on D. To avoid the nontrivial question of the exact PERSON feature geometry, I use descriptive label 3 for the relevant value ([−PARTICIPANT] as 3rd person). The syntactic structure for a pronoun referring to *doctor* is then as follows:

(21) $[_{DP}$ $D_{P:3,Num:\_\_}$ $[_{NumP/NP}$ ~~doctor~~$_{Num:sg}]]$

At the point of morphological realization, the morphology module cannot assign the right pronominal form until it determines whether PERSON corresponds to an animate referent. At the output of the syntax, the feature bundle gives no direct information about animacy. Whether or not the pronoun has an animate antecedent depends solely on the value assigned to the index. Vocabulary Insertion thus reflects the association of PERSON with an index by the CI interface.[38] If the index is animate, it becomes relevant whether the index corresponds to a female or male person.[39] The relevant assignment function determines whether the value of the index corresponds to an animate (human) referent. If it does, then the morphological realization needs to comply with Maximize Presupposition: Vocabulary Insertion realizes a feminine form if the referent is female and a masculine form if the referent is male or underspecified. A simplified derivation of *she$_{doctor}$* is given in (22).

(22) a. Syntactic output (after ϕ-feature matching and valuation)
         $[_{DP}$ $D_{P:3,Num:sg}$ $[_{NumP/NP}$ ~~doctor~~$_{Num:sg}]]$
      b. CI licensing and index association
         $[_{DP}$ $D_{P:3,Num:sg,i:7}$ $[_{NumP/NP}$ ~~doctor~~$_{Num:sg}]]$
      c. Vocabulary Insertion under assignment function $[7 \rightarrow Mary]$
         $D[3,\text{SG},i]\backslash\_\_\Delta_{NP}$ where i is a female person $\Rightarrow$ /she/

While this account suffices for English, more needs to be said about Italian as here a contextually valued gender feature may value an unvalued gender feature on agreeing elements, such as adjectival predicates. In Italian, unlike in English, D agrees in gender. It follows that D comes to the derivation with an unvalued gender feature. As for PERSON, it is introduced as a valued feature on D as in English. We can thus extend the hypothesis that PERSON is licensed and associated with *i* as part of semantically triggered labeling. Since a DP is an endocentric structure, the

---

[38] We could also stipulate that in English PERSON reflects animacy. As an anonymous reviewer points out, the English pronoun *it* could then be characterized as inanimate. An alternative is that *it* lacks a valued gender feature. Since a gender feature can be valued only in the presence of a person presupposition, we indirectly obtain the lack of animacy. I am not aware of any data that would distinguish between these two options for English. Since we would still need the gender-related step, adding animacy to the feature representation would not in and of itself resolve the primary question of gender marking.

The reviewer also raises the related question of the status of neuter in a three-way gender system. As far as I know, at least in Slavic, German, and Romanian, either neuter is introduced by the nominal representation in the lexicon, or it arises in the presence of a person presupposition or in defective-goal structures such as agreement of weather predicates or marking of nominalized infinitives. Although three-way systems would need to be investigated more carefully, this preliminary overview suggests that the present proposal can be extended to these gender systems as well.

[39] For quantifiers, this amounts to saying whether the quantificational domain is restricted to animate individuals of which the relevant properties are true.

features of the D head will become the label of the DP. In turn, the instantiation of PERSON on the D head itself becomes licensed and the unvalued gender feature becomes part of the label as well.

   As for the semantic valuation of the syntactically unvalued gender feature, I argue (following Preminger's (2009) logic) that ϕ-features may in principle remain unvalued. Yet they are reset to a—typically default—value as part of sending the syntactic output to the interfaces. I argue that if an unvalued gender feature cooccurs in the label with an animate PERSON feature associated with an index value, the gender feature is reset to a feminine gender value if the feminine gender presupposition is satisfied. The proposed implementation thus captures the observation that gender features may be generated as "free riders" on PERSON (e.g., Harley and Ritter 2002, Béjar and Rezac 2003, Adger and Harbour 2007).

   Let us turn to the rest of the derivation. Since features in the DP label are formally identical with their counterparts on D, the last-resort valuation of the gender feature in the label automatically values the gender feature on D as well. Since the gender feature on D has already established a matching Agree link with the unvalued gender feature on n, both instantiations of the gender feature are automatically valued. The morphology component then realizes the DP with the valued gender features. The derivation for a noun like *la chirurga* 'the.F surgeon.F' is schematized in (23).[40]
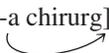
(23) la chirurga *'the.F surgeon.F'*
   a. Feature distribution from the lexicon and matching
      [$_{DP}$ D$_{[P:3,GEN:\_\_]}$] [$_n$ n$_{[GEN:\_\_]}$ √CHIRURG]]
   b. Labeling and CI licensing
      [$_{DP[P:3,GEN:\_\_,i:7]}$ D$_{[P:3,GEN:\_\_]}$ n . . .]
   c. Maximize Presupposition–driven gender valuation
      [$_{DP}$ D$_{[P:3,GEN:f,i:7]}$ D$_{[P:3,GEN:f]}$ n . . .]
   d. Automatic feature valuation and morphological output
      [$_{DP[P:3,GEN:f,i:7]}$ [$_{D[P:3,GEN:f]}$ la] [$_n$ [$_{n[GEN:f]}$ -a chirurg]]]

   Since the contextually driven feature valuation is a last-resort operation, it applies only if the gender feature is not valued in narrow syntax. If a gender feature is already associated with a noun from the lexicon, the gender feature on D obligatorily agrees with this lexically introduced value. Thus, the proposal captures the generalization that only a subset of nouns can have their value determined from the context. It also captures Matushansky's (2013) intuition that semantic gender features are a last-resort phenomenon.

   Crucially, valuation of the gender feature by the CI interface within a DP is possible only if the feature has not been previously valued from the lexicon.[41] That is, the gender feature on

---

[40] Strictly speaking, there should be no DP label in the first step of the derivation. I keep the traditional label in for clarity.
   [41] This follows from the timing of syntactic operations in (6).

D can be valued via PERSON only if the gender feature on D is unvalued at the point of Transfer. Thus, the valuation takes place in the configuration in (16) but not in the configuration in (17). Consequently, only a gender feature projected in a DP structure based on a lexical nominal representation of the liberal type (24a) or the semiconservative type (24b) can have a semantically interpretable gender feature. That is to say, only the gender feature of mating nouns (e.g., *il chirurg-o* 'the.M surgeon.M' vs. *la chirurg-a* 'the.F female surgeon.F') and semiconservative nouns (i.e., noun forms based on a class feature but for which agreeing elements in the DP are determined by the natural gender of the referent; e.g., *il chirurg-o* 'the.M surgeon' vs. *la chirurg-o* 'the.F female surgeon') is semantically interpretable.

(24) *Lexicon representations of nouns compatible with natural gender*
   a. Liberal (*il chirurgo*, *la chirurga*): √root
   b. Semiconservative (*il chirurgo*, *la chirurgo*): [$_n$ n$_{[CL:1]}$ √root]

The consequence is that only a gender feature derived from PERSON is semantically interpreted. Grammatical gender introduced in the lexicon is not.[42] No distinction between interpretable and uninterpretable syntactic features is needed to derive the interpretive distinction. Furthermore, the proposed implementation sheds light on the descriptive distinction between so-called animate and inanimate gender. Technically, animate gender is a gender feature that requires the cooccurrence of PERSON with an animate referent. Inanimate gender lacks this cooccurrence requirement.

2.3 More on the Role of Timing: Spell-Out vs. Transfer

A core property of the syntax-semantics system proposed here is that there is a point in the derivation when the narrow-syntax representation can interact with the semantic component: namely, values licensed at the CI interface are accessible to narrow syntax only when a phase is completed. In this section, I argue that while there is only one point in the derivation when syntax features can be valued by the CI interface, there are *two* distinct points when the structure can be sent to the morphology. The derivational distinction corresponds to the independently proposed distinction between the Spell-Out and the Transfer of a phase. I will utilize this distinction and argue that it plays a substantial role in the Italian gender system.

I model Spell-Out as the point in the derivation when narrow-syntax operations are completed. By Transfer, I mean the point in the derivation when a phase is sent to the CI interface.[43] Once the structure is spelled out, no new syntactic relations can be created. Here the independently proposed distinction between matching and valuation as components of Agree becomes crucial. While narrow syntax cannot create a new matching link after Spell-Out, an existing matching

---

[42] See also Bobaljik and Zocca 2011 for a related set of facts.
[43] An anonymous reviewer raises the important question of how the proposed model relates to the notion of LF as in Heim and Kratzer 1998. Although this might turn out to be a vast simplification, there is no fundamental difference between LF and the model of the CI interface used here. Under this view, Transfer maps the narrow-syntax structure onto a representation that may be further modified at the CI interface/LF. For instance, some syncategorematic elements sensitive to the narrow-syntax properties of their environment, such as a λ-abstractor introduced for some instances of movement (Ā-movement but perhaps not all instances of A-movement), could be part of the Transfer procedure.

link can be valued at the interfaces.[44] In the case investigated here, this means that an unvalued matched feature can be valued by the CI interface. After the feature is valued and the structure is transferred, such a semantically enriched syntactic representation is realized by the morphology module. I argue that an economy condition that prefers the two interfaces to be isomorphic (be it Heim's (1991) Maximize Presupposition or an economy condition of the sort proposed in Williams 2003, Reinhart 2006, Bobaljik and Wurmbrand 2012, and Kučerová 2012, among others) can apply only during the temporal window between Spell-Out and Transfer.

In turn, the realized morphological representation can reflect semantic features that were not originally present in the narrow-syntax derivation only if the phase was sent to the morphology module at the point of Transfer, instead of Spell-Out. A model of the proposed grammar architecture is shown in (25).

(25) *Model of a derivation with Transfer preceding morphology*

Morphology   CI

Transfer ~ CI valuation

Spell-Out

Syntax

Lexicon
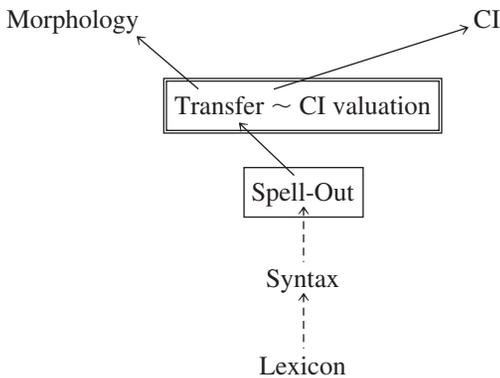
This derivational trajectory accounts for Italian nouns in which either the form of the noun itself or the agreeing elements within the extended DP morphologically reflect the natural gender of their referent: specifically, mating nouns, which I argue are based on liberal roots, such as *il bambino*, *la bambina* and *il chirurgo*, *la chirurga*, and semiconservative nouns, which I argue come from the lexicon with a class marker but without a valued gender feature, such as *il chirurgo*, *la chirurgo*. In this type of structure, there is no valued gender within the DP in narrow syntax. The semantically informed value is assigned when the DP phase is transferred. Since the morphology realizes the phase only after it undergoes Transfer, the morphology faithfully reflects the semantically valued gender feature. A detailed derivation of this type was given in (23).

The empirical question is whether a syntactic derivation could be sent to the morphology *before* the phase is transferred and the features are valued by the CI interface. I argue that an Italian speaker's grammar indeed allows for both derivational options: the phase can be sent to the morphology *before* or *after* it undergoes Transfer. If the phase is sent to the morphology

---

[44] This amounts to saying that narrow syntax requires unvalued features to be matched but not necessarily to be valued. See Preminger 2009 for empirical arguments in favor of this view.

before it undergoes Transfer, the morphological realization of the DP cannot realize a semantically valued gender. The corresponding derivational trajectory is schematized in (26).

(26) *Model of a derivation with Transfer parallel to morphology*

Morphology   CI

Transfer ~ CI valuation

Spell-Out

Syntax
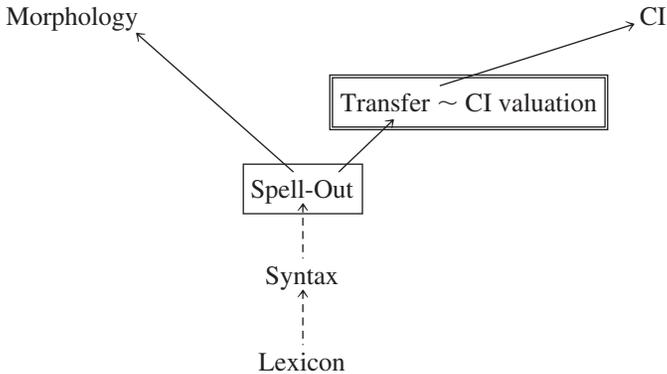
Lexicon

    The empirical support for this timing distinction comes from the data discussed in section 1.3. We saw that lexically semiconservative nouns exhibit two distinct agreement patterns. That is, if a noun like *soprano* comes from the lexicon with a class specification and if it denotes a female referent, the form of the noun is invariably *sopran-o* and it triggers feminine agreement on predicates. However, the elements *within* the DP agree either in the feminine, as in (27a), or in the masculine, as in (27b).

(27) a. La   soprano è andata.
       the.[F] soprano is gone.[F]
    b. Il   soprano è andata.
       the.[M] soprano is gone.[F]

    The former pattern, (27a), corresponds to the derivation schematized in (25): that is, if the DP-internal and DP-external agreements match, the narrow-syntax representation of the DP was enriched by the CI interface before the DP was sent to the morphology module. In other words, the morphological realization of the DP reflects the label of the DP.

    The surprising pattern is (27b), where the DP-internal elements appear as masculine. I argue that this agreement pattern arises only if the narrow-syntax representation of the DP is sent to the morphology immediately after Spell-Out, that is, before the phase is transferred and the DP is labeled. Since the syntactic structure contains an unvalued ϕ-feature, the morphology must resort to a default morphological realization, which in Italian corresponds to masculine. In other words, if the internal agreement is masculine, it is an effect of *failed Agree* in Preminger's (2009) sense. Crucially, the failed-Agree realization arises only within the morphology. That is, the narrow-syntax feature valuation remains unaltered.

    However, for the DP to be selected for further syntactic derivation (i.e., to be merged with other syntactic material), it must first undergo Transfer and be labeled in the process. At this point, the unvalued gender feature can be valued by the CI interface. Although this feature valuation is not morphologically reflected in the DP itself, as the structure has already been spelled out, the

feature valued by the CI interface is available for further Agree relations, such as Agree between the transferred DP and a predicate. The reason is that Agree targets features in the label, and the semantically valued gender is part of the label. Thus, the feminine agreement on the participle in (27b) is a morphological reflex of syntactic Agree between the unvalued gender feature on the participle and the valued gender feature in the label of the transferred DP.

A crucial property of the proposed system is that a feature can be valued by the CI interface only as a last resort. The proposal thus makes a clear prediction: if a valued feature is present in the structure (introduced from the lexicon), the label of the DP cannot reflect a semantically enriched value of the feature, and in turn no local Agree link can be valued by a semantically valued feature. This is precisely what happens with conservative nouns. For speakers who store *chirurgo* in the lexicon with a class feature and a masculine gender feature, the masculine gender feature values the unvalued gender feature on D and in turn becomes part of the DP label. Since the gender feature in the label is valued as masculine, all local Agree links are valued as masculine even if the intended label denotes a woman, as in (28). The corresponding derivation is schematized in (29). In other words, once the label is set, the values cannot be altered. Only unvalued features in the label can be valued as a last resort, to comply with Maximize Presupposition.

(28) *Lexically given gender in the label overrides Maximize Presupposition*
     Il   chirurgo è andato.
     the.[M] surgeon is gone.[M]

(29) il chirurgo *'the.M surgeon.M': Conservative structure, denoting a female*
    a. Feature distribution from the lexicon and matching
       $[_{DP} D_{[P:3,GEN:\_\_]} [_n n_{[GEN:m,CLASS:1]} \sqrt{CHIRURG}]]$

    b. Feature valuation from the structure
       $[_{DP} D_{[P:3,GEN:m]} [_n n_{[GEN:m,CLASS:1]} \sqrt{CHIRURG}]]$

    c. Labeling and licensing by the CI interface (where [7 → Mary])
       $[_{DP[P:3,GEN:m,i:7]} D_{[P:3,GEN:m]} n \ldots]$

    d. No effect of Maximize Presupposition and morphological output
       $[_{DP[P:3,GEN:m,i:7]} [_{D[P:3,GEN:m]} il] [_n [_{n[GEN:m]} -o] chirurg]]$

The proposal predicts that if a noun comes from the lexicon with a valued gender feature, any local agreement must respect the lexically given gender specification. That is to say, if a noun comes from the lexicon with a valued gender feature that does not match the natural gender of the noun's referent, the natural gender can be reflected only in nonlocal agreement (e.g., across a sentence boundary); it can never be reflected in a local agreement configuration, where by *local* I mean an Agree relation established in narrow syntax.[45] To explore this prediction, we need to look at grammatically feminine nouns in order to control for morphological default realizations.

---

[45] The predictions of the present model thus differ radically from those of models that assume the presence of a pronoun within D or of two distinct types of features, as in Collins and Postal 2012.

Grammatically feminine nouns like *guida* 'guide' and *guardia* 'guard' obligatorily trigger feminine agreement on predicates, regardless of the natural gender of their referent. Ferrari-Bridgers (2007) observes, however, that there is a striking exception to the pattern: if such a noun appears in a nominal conjunction, then it triggers context-dependent agreement. While for Ferrari-Bridgers the pattern means that this type of noun is underspecified for gender, I argue that it provides evidence for the grammar architecture proposed here. Since such a noun comes from the lexicon with valued gender, natural gender agreement is restricted to nonlocal syntactic operations. The switch is demonstrated in (30). The participle in (30a) is valued by Agree with the DP that was labeled by the lexically specified gender. In other words, local agreement must reflect the lexically specified value. In contrast, if such a DP is embedded in a coordination, participle agreement with the coordination reflects the natural gender of the referents, as in (30b–c).

(30) a. La brava guida   si   e'persa nel   bosco.
        the good guide.[F] him/her lost.[F] in.the woods
        'The good guide lost his/her way in the forest.'
     b. La guardia e   sua sorella sono andate   al   cinema sta sera.
        the guard.[F] and self sister   have gone.[F.PL] to.the movies this evening
        'The guard and her sister went to the movies tonight.'
     c. La guardia e   sua sorella son andati   al   cinema sta sera.
        the guard.[F] and self sister   have gone.[M.PL] to.the movies this evening
        'The guard and his sister went to the movies tonight.'
        (adapted from Ferrari-Bridgers 2007:151,(4))

I argue that the difference in the agreement patterns reflects their distinct locality properties. That is, in the case of agreement with a coordination, Agree cannot directly access the DP in question; instead, the participle probes the label of the coordination. Since conjunction forms a plurality (Munn 1993, Bošković 2009, Bhatt and Walkow 2013) and this is a process that requires access to the semantic component, the compounded gender value of the two conjuncts is based on the value provided by the CI interface, instead of being determined by the lexicon. In turn, the natural gender value is accessible to agreement only if Agree targets a semantically specified label, instead of the φ-features present in the label of the DP.[46]

Another prediction concerns a crosslinguistic variation. The current proposal assumes that within a DP, masculine behaves like a semantically unmarked feature because the masculine morphology arises via failed Agree; that is, there is no valued syntactic feature to be interpreted by the semantics. If a language resorted to feminine as its default failed-Agree value, we would expect feminine to be associated with a wider variety of semantic interpretation. Crosslinguistic

---

[46] An anonymous reviewer asks whether coordination could be derived by a morphosyntactic mechanism, instead of referring to semantic plurality. I am not aware of a proposal that would successfully account for the difference between *a good friend and his editor was/were* . . . without using some semantic component. Thus, I follow proposals that argue for some combination of semantic and morphosyntactic valuation (e.g., Farkas and Zec 1995, King and Dalrymple 2004, Heycock and Zamparelli 2005).

investigation of this prediction goes beyond the scope of this article, but some Arabic dialects, for instance, might be good candidates for testing it.

2.4 Further Extensions

Although the proposed grammar architecture model was specifically developed for Italian nominal inflection, if it is on the right track, it should extend to other empirical domains as well. In general, the model predicts that the grammatical and semantic values of φ-features are not necessarily established at the same point in the derivation; consequently, derivational timing differences correspond to interpretive differences. We thus expect to find a parallel derivational effect in other languages and with other features. I briefly discuss three sets of data that provide suggestive evidence for the proposed model.

Since the foundational work of Picallo (1991) and Ritter (1993), it has been known that gender valuation varies crosslinguistically. As Ritter observes, nominal derivational morphology in Hebrew does not preserve the gender of the root noun. However, the gender of a Hebrew noun cannot switch because of the natural gender of its referent. The opposite is true of Spanish (Harris 1991). In the context of this article, neither of these facts is surprising. Natural-gender-dependent switching (mating nouns) is possible only if a noun comes from the lexicon without a lexically valued gender feature.[47] While this is occasionally possible in Romance languages, Hebrew does not seem to have this option.[48] Similarly, if the form of the root noun is in principle independent of the lexically determined gender feature value, it is not surprising that in languages like Hebrew, a word that has the phonotactic shape of a masculine noun can still trigger feminine agreement, as shown in (31).

(31) a. šana   tova
        year.F.SG good.F.SG
      *šana   tov
        year.F.SG good.M.SG
   b. šanim   tovot
        year.F.PL good.F.PL
      *šanim   tovim
        year.F.PL good.M.PL
        (Ritter 1993:799, (8))

Related data come from German. As Yatsushiro and Sauerland (2006) observe, there are instances of morphologically marked gender features that may but do not have to be composition-

---

[47] That nouns can differ in whether or not their gender feature is lexically specified has been independently argued by Alexiadou (2004) for Romance, Greek, and Hebrew and by Riente (2003) for Italian.

[48] Languages differ in how nominal derivational morphology treats gender. Languages such as German resemble Hebrew in that the gender of a diminutive is solely determined by the derivational morphology (diminutives are always neuter in German). However, in other languages, such as Russian, the gender of a diminutive tends to retain the lexically determined value of the nominal root. See, for example, Wiltschko and Steriopolo 2007 and Matushansky 2015.

ally interpreted.[49] The critical distinction can be seen in superlative DPs that involve the derivational morpheme *-in*, which forms profession nouns for females in German, as in (32).

(32) Merkel jetzt beliebteste   Politiker-in Deutschlands   (headline)
     Merkel now most.popular politician-F of.Germany
     a. 'Merkel now most popular female politician in Germany' [= true even if there is
        a male politician more popular than Merkel]
     b. 'Merkel now most popular politician in Germany' [= false if there is a male politi-
        cian more popular than Merkel]
        (Yatsushiro and Sauerland 2006:11, (2))

If the profession noun *Politiker* combines with the feminine ending *-in*, the superlative meaning of the complex DP either can be calculated with respect to female politicians only (then *-in* is semantically active and the superlative morpheme scopes over it) or can refer to politicians of both genders (then *-in* is not semantically interpreted; instead, it is a pure "agreement" reflex of the natural gender associated with the individual Merkel, which can be modeled as being in the scope of the superlative morpheme). In a theory that associates interpretability of gender features with a distinct height of functional heads, either two distinct derivational morphemes would be required, or a single morpheme would need to have two distinct functional projections for gender.

   Neither option seems well-motivated. The problem is simpler in the present account and can be straightforwardly tied to the Spell-Out/Transfer timing distinction. When the morpheme *-in* is morphologically realized before the CI interface can associate the gender feature with a semantic index corresponding to a feminine referent and in turn activate the gender presupposition that would enforce the female interpretation, we obtain reading (32a). When the feature bundle is sent to the morphology after the semantic index is added, we obtain reading (32b).

   An analogical distinction in derivational timing can be shown for number features as well. In some British English dialects, group nouns like *team* can trigger both singular and plural agreement. The seeming optionality disappears in certain contexts: if such a noun is an indefinite, the agreement pattern restricts the scope of the indefinite. Consider (33).

(33) A northern team is/are certain to be in the final.
     a. is: ∃ > *certain*, *certain* > ∃
     b. are: ∃ > *certain*, *certain > ∃
     (Sauerland and Elbourne 2002:288, (14))

While singular agreement is compatible with two distinct LF interpretations, plural agreement forces a wide scope of the indefinite.[50] Unless we posit a lookahead derivation, this correlation

---

[49] For an early and comprehensive argument that morphological features corresponding to functional categories do not always combine compositionally, see Carlson 1983.

[50] That is to say, for the plural agreement proposition to be true, there must be a particular team that is certain to be in the final. In contrast, for the singular agreement proposition to be true, there may be either a particular team from the north or only teams from the north that is/are certain to be in the final (in which case, a northern team is certain to be in the final).

between agreement and a particular LF representation strongly suggests that the grammatically based singular agreement is established *before* the corresponding LF representation is determined. In contrast, the semantically informed plural agreement strictly refers to a point in the derivation *after* the potential scope ambiguity has been resolved.[51] Theories that attribute differences in grammatical and semantic features to distinct functional heads within an extended DP have little to say about this difference in derivational timing. In contrast, the theory proposed here has a straightforward answer for the observed asymmetry. If the singular (i.e., grammatical) agreement is established before Transfer, it does not reflect the semantic interpretation. In contrast, if the plural agreement is established only as part of Transfer, it reflects the semantic representation.[52] This of course does not account for all the properties of the pattern, that is, why only wide scope triggers the semantic agreement and why semantic agreement arises at all if there is a valued grammatical feature in the structure. The overall pattern, however, is compatible with the model of ϕ-feature valuation introduced in this article.

3 Conclusions

This article provides an explicit model of Standard Italian nominal inflection throughout the course of its derivation: first its representation in the lexicon, then its derivation in narrow syntax, and finally its realization at the morphological and semantic interfaces. However, extensions of the proposed model—especially the proposed model of ϕ-features at the syntax-semantics interface—go beyond Italian.

The article makes two major theoretical contributions. The first, concerning lexical representation of nominal structures, advances our understanding of root representations in the lexicon and the distinction between interpretable and uninterpretable features. The article provides evidence that nouns can come from the lexicon with or without a valued gender feature. An unvalued gender feature may be valued by the CI interface when the DP is sent to the syntax-semantics interface. The logic of the proposal is that a lexically determined gender feature is never semantically interpreted, a fact often masked by the lexical semantics of the corresponding root. A gender feature is interpreted only if it is valued at the syntax-semantics interface. That is, the distinction between interpretable and uninterpretable features is reduced to independently needed properties of the assumed grammar architecture. Second, the article proposes a theory of ϕ-features at the syntax-semantics interface that relies strictly on the Minimalist Program grammar architecture. The proposed model provides evidence that there are two distinct routes for a syntactic representation to

---

[51] The present account is very different from Sauerland and Elbourne's (2002). I do not follow their analysis because it incorrectly predicts that grammatically plural nouns that denote singular objects, such as pluralia tantum, should optionally trigger singular agreement. As far as I can tell, the symmetry problem is inherent to their account and cannot be fixed without additional stipulations.

[52] As an anonymous reviewer points out, this conclusion runs counter to the observation that plural marking in English corresponds to an interpretive default for number (Sauerland 2003, Sauerland, Andersen, and Yatsushiro 2005). Yet, if the number value is contextually provided, the value must match the semantic value (in the sense of plural as mereological or Boolean conjunction). I don't have a conclusive answer about when plural marking corresponds to a semantic default. However, if the analysis proposed here for the interpretation of gender is on the right track, association of the syntactic number feature with a semantic index could be a key element.

be sent to the morphology: either the narrow-syntax representation is sent to the morphology at the point of Spell-Out and then the morphological realization reflects the narrow-syntax feature configuration, or it is sent to the morphology at Transfer—that is, with the syntactic representation enriched by semantic information. The latter route ties the Italian data patterns to data that have been independently argued to involve a derivational asymmetry whereby morphological realization is sensitive to semantic representation (e.g., Williams 2003, Reinhart 2006, Bobaljik and Wurmbrand 2012, Kučerová 2012).

A core property of the syntactic implementation is that although the locus of valuation of the gender feature may differ, grammatically both types of gender feature—grammatical and natural gender—are realized within the same Agree chain. This property, combined with the two possible derivational routes, accounts for another hitherto unexplained property of the grammatical vs. semantic feature distinction. If there are independent instances of gender feature valuation—a proposal spelled out most precisely by Kramer (2009, 2015)—we would expect to find gender stacking, analogical to number and case stacking (Hagit Borer, pers. comm.). This prediction is not borne out in the current analysis, suggesting that it is incorrect. Indeed, I am not aware of any language that has such a property. Even though there may be more than one source of gender feature valuation, there is only one Agree chain within an extended DP. Thus, this structural property accounts for the lack of gender stacking.

In addition to providing a theoretical motivation for distinguishing between syntactic valuation and contextual valuation as a direct consequence of the phase-based grammar architecture, the article makes two further theoretical contributions. The first is that only a subset of nouns are based on category-neutral roots (cf. Marantz 1997), a conclusion supported by data from derivational morphology. Second, the proposal sheds greater light on the crosslinguistically pervasive generalization that animate 3rd person nominals share feature properties with 1st and 2nd person nominals (e.g., Ormazabal and Romero 1998, 2007, Adger and Harbour 2007, Ritter 2014, Welch 2014, Lochbihler and Oxford 2015, Wiltschko and Ritter 2015). This follows under the hypothesis that animate nouns share the [+participant] feature with 1st and 2nd person nominals, and are licensed by the CI interface in the same manner. Since 3rd person inanimate nominals have a different person feature valuation, they may be treated differently at the CI interface.

Finally, the article opens the door to rethinking a crosslinguistically prevalent dependency of gender on number. If semantically interpretable gender is indeed dependent on PERSON and there is a connection between semantic individuation and PERSON, there might be no direct dependency between gender and number (see Greenberg 1963, Noyer 1992, Harley and Ritter 2002). Fully exploring this hypothesis is a topic for future research.

References

Acquaviva, Paolo. 2008. *Lexical plurals: A morphosemantic approach.* Oxford: Oxford University Press.

Adger, David. 2003. *Core syntax: A Minimalist approach.* Oxford: Oxford University Press.

Adger, David, and Daniel Harbour. 2007. Syntax and syncretisms of the Person Case Constraint. *Syntax* 10:2–37.

Alexiadou, Artemis. 2004. Inflection class, gender and DP internal structure. In *Explorations in nominal inflection*, ed. by Gereon Müller, Lutz Gunkel, and Gisela Zifonun, 21–50. Berlin: Mouton de Gruyter.

Alexiadou, Artemis. 2011. Plural mass nouns and the morpho-syntax of number. In *Proceedings of the 28th West Coast Conference on Formal Linguistics*, ed. by Mary Byram Washburn, Katherine McKinney-Bock, Erika Varis, Ann Sawyer, and Barbara Tomaszewicz, 33–41. Somerville, MA: Cascadilla Proceedings Project.

Arsenijević, Boban, and Wolfram Hinzen. 2012. On the absence of X-within-X recursion in human grammar. *Linguistic Inquiry* 43:423–440.

Béjar, Susana, and Milan Rezac. 2003. Person licensing and the derivation of PCC effects. *Amsterdam Studies in the Theory and History of Linguistic Science Series* 4:49–62.

Bhatt, Rajesh, and Martin Walkow. 2013. Locating agreement in grammar: An argument from agreement in conjunctions. *Natural Language and Linguistic Theory* 31:951–1013.

Bobaljik, Jonathan David, and Susi Wurmbrand. 2012. Word order and scope: Transparent interfaces and the 3/4 signature. *Linguistic Inquiry* 43:371–421.

Bobaljik, Jonathan David, and Cynthia Levart Zocca. 2011. Gender markedness: The anatomy of a counterexample. *Morphology* 21:141–166.

Borer, Hagit. 2005. *Structuring sense.* Vol. 1, *In name only.* Oxford: Oxford University Press.

Bošković, Željko. 2009. Unifying first and last conjunct agreement. *Natural Language and Linguistic Theory* 27:455–496.

Carlson, Greg N. 1983. Marking constituents. In *Linguistic categories: Auxiliaries and related puzzles*, ed. by Frank Heny and Barry Richards, 1:69–98. Dordrecht: Reidel.

Chomsky, Noam. 2000. Minimalist inquiries: The framework. In *Step by step*, ed. by Roger Martin, David Michaels, and Juan Uriagereka, 89–155. Cambridge, MA: MIT Press.

Chomsky, Noam. 2001. Derivation by phase. In *Ken Hale: A life in language*, ed. by Michael Kenstowicz, 1–52. Cambridge, MA: MIT Press.

Chomsky, Noam. 2008. On phases. In *Foundational issues in linguistic theory: Essays in honor of Jean-Roger Vergnaud*, ed. by Robert Freidin, Carlos P. Otero, and Maria Luisa Zubizarreta, 133–166. Cambridge, MA: MIT Press.

Chomsky, Noam. 2013. Problems of projection. *Lingua* 130:33–49.

Chomsky, Noam. 2015. Problems of projection: Extensions. In *Structures, strategies and beyond: Studies in honour of Adriana Belletti*, ed. by Elisa Di Domenico, Cornelia Hamann, and Simona Matteini, 1–16. Amsterdam: John Benjamins.

Collins, Chris, and Paul M. Postal. 2012. *Imposters: A study of pronominal agreement.* Cambridge, MA: MIT Press.

Corbett, Greville. 1991. *Gender.* Cambridge: Cambridge University Press.

Dahl, Östen. 2000. Animacy and the notion of semantic gender. In *Gender in grammar and cognition.* Vol. 1, *Approaches to gender*, ed. by Barbara Unterbeck and Matti Rissanen, 99–115. Berlin: Mouton de Gruyter.

Danon, Gabi. 2011. Agreement and DP-internal feature distribution. *Syntax* 14:297–317.

Dressler, Wolfgang U., and Anna M. Thornton. 1996. Italian nominal inflection. *Wiener linguistische Gazette* 57:1–26.

Elbourne, Paul. 2005. *Situations and individuals.* Cambridge, MA: MIT Press.

Embick, David. 2000. Features, syntax, and categories in the Latin perfect. *Linguistic Inquiry* 31:185–230.

Embick, David, and Morris Halle. 2005. On the status of stems in morphological theory. In *Romance Languages and Linguistic Theory 2003: Selected papers from 'Going Romance' 2003, Nijmegen, 20–22 November*, ed. by Twan Geerts, Ivo van Ginneken, and Haike Jacobs, 37–72. Amsterdam: John Benjamins.

Embick, David, and Rolf Noyer. 2007. Distributed Morphology and the syntax-morphology interface. In *The Oxford handbook of linguistic interfaces*, ed. by Gillian Ramchand and Kurt Rice, 289–324. Oxford; Oxford University Press.

Endriss, Cornelia. 2009. *Quantificational topics: A scopal treatment of exceptional wide scope phenomena.* New York: Springer.

Fabrizio, Claudia. 2013. The meaning of a noun converted into a verb: A semantic exploration of Italian. *Rivista di Linguistica* 25:175–219.

Farkas, Donka, and Draga Zec. 1995. Agreement and pronominal reference. In *Advances in Roumanian linguistics*, ed. by Guglielmo Cinque and Giuliana Giusti, 83–102. Amsterdam: John Benjamins.

Ferrari, Franca. 2005. A syntactic analysis of the Italian and Luganda nominal systems: How nouns can be formed in the syntax. Doctoral dissertation, New York University.

Ferrari-Bridgers, Franca. 2007. The predictability of gender in Italian. *Lingua et Linguistica* 1:146–167.

Ferrari-Bridgers, Franca. 2008. A unified syntactic analysis of Italian and Luganda nouns. In *The Bantu-Romance connection: A comparative investigation of verbal agreement, DPs, and information structure*, ed. by Cécile De Cat and Katherine Demuth, 239–260. Amsterdam: John Benjamins.

Greenberg, Joseph H. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In *Universals of language*, ed. by Joseph H. Greenberg, 73–113. Cambridge, MA: MIT Press.

Halle, Morris, and Alec Marantz. 1993. Distributive Morphology and the pieces of inflection. In *The view from Building 20: Essays in linguistics in honor of Sylvain Bromberger*, 111–176. Cambridge, MA: MIT Press.

Harley, Heidi, and Elizabeth Ritter. 2002. Person and number in pronouns: A feature-geometric analysis. *Language* 78:482–526.

Harris, James W. 1991. The exponence of gender in Spanish. *Linguistic Inquiry* 22:27–62.

Heim, Irene. 1991. Artikel und Definitheit. In *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, ed. by Arnim von Stechow and Dieter Wunderlich, 487–535. Berlin: Mouton de Gruyter.

Heim, Irene. 2008. Features on bound pronouns. In *Phi-theory: Phi features across interfaces and modules*, ed. by Daniel Harbour, David Adger, and Susana Béjar, 35–56. Oxford: Oxford University Press.

Heim, Irene, and Angelika Kratzer. 1998. *Semantics in generative grammar.* Oxford: Blackwell.

Heycock, Caroline, and Roberto Zamparelli. 2005. Friends and colleagues: Plurality, coordination, and the structure of DP. *Natural Language Semantics* 13:201–270.

King, Tracy Holloway, and Mary Dalrymple. 2004. Determiner agreement and noun conjunction. *Journal of Linguistics* 40:69–104.

Kramer, Ruth. 2009. Definite markers, phi-features, and agreement: A morphosyntactic investigation of the Amharic DP. Doctoral dissertation, University of California, Santa Cruz.

Kramer, Ruth. 2015. *The morphosyntax of gender.* Oxford: Oxford University Press.

Kratzer, Angelika. 2009. Making a pronoun: Fake indexicals as windows into the properties of pronouns. *Linguistic Inquiry* 40:187–237.

Kučerová, Ivona. 2012. Grammatical marking of givenness. *Natural Language Semantics* 20:1–30.

Lampitelli, Nicola. 2010. Nounness, gender, class and syntactic structures in Italian nouns. In *Romance Languages and Linguistic Theory 2008: Selected papers from 'Going Romance' Groningen 2008*, ed. by Reineke Bok-Bennema, Brigitte Kampers-Manhe, and Bart Hollebrandse, 195–214. Amsterdam: John Benjamins.

Lampitelli, Nicola. 2014. The Romance plural isogloss and linguistic change: A comparative study of Romance nouns. *Lingua* 140:158–179.

Landau, Idan. 2010. *The locative syntax of experiencers.* Cambridge, MA: MIT Press.

Landau, Idan. 2016. DP-internal semantic agreement: A configurational analysis. *Natural Language and Linguistic Theory* 34:975–1020.

Larson, Richard K. 2011. Clauses, propositions, and phases. In *The biolinguistic enterprise: New perspectives on the evolution and nature of the human language faculty*, ed. by Anna Maria Di Sciullo and Cedric Boeckx, 366–391. Oxford: Oxford University Press.

Lochbihler, Bethany, and Will Oxford. 2015. The person-animacy connection in Algonquian. Paper presented at the 2nd Prairies Workshop on Language and Linguistics. University of Manitoba, Winnipeg, 14 March 2015.

Longobardi, Giuseppe. 2008. Reference to individuals, person, and the variety of mapping parameters. In *Essays on nominal determination: From morphology to discourse management*, ed. by Henrik Høeg Müller and Alex Klinge, 189–211. Amsterdam: John Benjamins.

Lowenstamm, Jean. 2008. On little n, √, and types of nouns. In *Sounds of silence: Empty elements in syntax and phonology*, ed. by Jutta M. Hartmann, Veronika Hegedűs, and Henk van Riemsdijk, 106–143. Amsterdam: Elsevier.

Marantz, Alec. 1997. No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. In *Proceedings of the 21st Annual Penn Linguistics Colloquium*, ed. by Alexis Dimitriadis, Laura Siegel, Clarissa Surek-Clark, and Alexander Williams, 201–225. Pennsylvania Working Papers in Linguistics 4.2. Philadelphia: University of Pennsylvania, Penn Linguistics Club.

Matushansky, Ora. 2013. Gender confusion. In *Diagnosing syntax*, ed. by Lisa Lai-Shen Cheng and Norbert Corver, 271–294. Oxford: Oxford University Press.

Matushansky, Ora. 2015. N is for "not there." Paper presented at the FASL 24 Workshop on Approaches to Slavic Morphology, New York University, 7 May 2015.

Merchant, Jason. 2014. Gender mismatches under nominal ellipsis. *Lingua* 151:9–32.

Munn, Alan. 1993. Topics in the syntax and semantics of coordinate structures. Doctoral dissertation, University of Maryland, College Park.

Narita, Hiroki. 2011. Phasing in full interpretation. Doctoral dissertation, Harvard University, Cambridge, MA.

Neeleman, Ad. 2008. Subset controllers. Paper presented at Words Don't Come Easy, University of Verona, 17–18 November 2008.

Nevins, Andrew. 2007. The representation of third person and its consequences for person-case effects. *Natural Language and Linguistic Theory* 25:273–313.

Norris, Mark. 2014. A theory of nominal concord. Doctoral dissertation, University of California, Santa Cruz.

Noyer, Rolf. 1992. Features, positions and affixes in autonomous morphological structure. Doctoral dissertation, MIT, Cambridge, MA.

Oltra-Massuet, Isabel. 1999. On the notion of theme vowel: A new approach to Catalan verbal morphology. Doctoral dissertation, MIT, Cambridge, MA.

Ormazabal, Javier, and Juan Romero. 1998. On the syntactic nature of the *me-lui* and the Person-Case Constraint. *Anuario del Seminario Julio de Urquijo* 32:415–434.

Ormazabal, Javier, and Juan Romero. 2007. The Object Agreement Constraint. *Natural Language and Linguistic Theory* 25:315–347.

Ouwayda, Sarah. 2014. Where number lies: Plural marking, numerals, and the collective-distributive distinction. Doctoral dissertation, University of Southern California, Los Angeles.

Percus, Orin. 2011. Gender features and interpretation: A case study. *Morphology* 21:167–196.

Pesetsky, David. 2013. *Russian case morphology and the syntactic categories.* Cambridge, MA: MIT Press.

Pesetsky, David, and Esther Torrego. 2007. The syntax of valuation and the interpretability of features. In *Phrasal and clausal architecture: Syntactic derivation and interpretation. In honor of Joseph E. Emonds*, ed. by Simin Karimi, Vida Samiian, and Wendy K. Wilkins, 262–294. Amsterdam: John Benjamins.

Pesetsky, David, and Esther Torrego. 2011. Case. In *The Oxford handbook of linguistic Minimalism*, ed. by Cedric Boeckx, 52–72. Oxford: Oxford University Press.

Picallo, M. Carme. 1991. Nominals and nominalizations in Catalan. *Probus* 3:279–316.

Postal, Paul M. 1969. On so-called 'pronouns' in English. In *Modern studies in English*, ed. by David Reibel and Sanford A. Schane, 201–244. Englewood Cliffs, NJ: Prentice-Hall.

Preminger, Omer. 2009. Breaking agreements: Distinguishing agreement and clitic doubling by their failures. *Linguistic Inquiry* 40:619–666.

Reinhart, Tanya. 2006. *Interface strategies: Optimal and costly computations.* Cambridge, MA: MIT Press.

Riente, Lara. 2003. Ladies first: The pivotal role of gender in the Italian nominal inflection system. *McGill Working Papers in Linguistics* 17(2):1–54.

Ritter, Elizabeth. 1993. Where's gender? *Linguistic Inquiry* 24:795–803.

Ritter, Elizabeth. 1995. On the syntactic category of pronouns and agreement. *Natural Language and Linguistic Theory* 13:405–443.

Ritter, Elizabeth. 2014. Featuring animacy. *Nordlyd* 41:103–124.

Sauerland, Uli. 2002. The present tense is vacuous. *Snippets* 6:12–13.

Sauerland, Uli. 2003. A new semantics for number. In *Proceedings of SALT 13*, ed. by Robert B. Young and Yuping Zhou, 258–275. http://journals.linguisticsociety.org/proceedings/index.php/SALT/issue /view/96.

Sauerland, Uli. 2004. A comprehensive semantics for agreement. Paper presented at Phi-Workshop, McGill University, Montreal.

Sauerland, Uli. 2008. Implicated presuppositions. In *The discourse potential of underspecified structures*, ed. by Anita Steube, 581–600. Berlin: Mouton de Gruyter.

Sauerland, Uli, Jan Andersen, and Kazuko Yatsushiro. 2005. The plural is semantically unmarked. In *Linguistic evidence: Empirical, theoretical, and computational perspectives*, ed. by Stephan Kepser and Marga Reis, 413–434. Berlin: Mouton de Gruyter.

Sauerland, Uli, and Paul Elbourne. 2002. Total reconstruction, PF movement, and derivational order. *Linguistic Inquiry* 33:283–319.

Smith, Peter. 2015. Feature mismatches: Consequences for syntax, morphology and semantics. Doctoral dissertation, University of Connecticut, Storrs.

Spathas, Giorgos. 2010. *Focus on anaphora.* Utrecht: LOT.

Stechow, Arnim von. 2003. Feature deletion under semantic binding: Tense, person, and mood under verbal quantifiers. In *NELS 33*, ed. by Makoto Kadowaki and Shigeto Kawahara, 379–404. Amherst: University of Massachusetts, Graduate Linguistic Student Association.

Steriopolo, Olga, and Martina Wiltschko. 2010. Distributed Gender hypothesis. In *Formal Studies in Slavic Linguistics: Proceedings of the Formal Description of Slavic Languages 7.5*, ed. by Gerhild Zybatow, Philip Dudchuk, Serge Minor, and Ekaterina Pshehotskaya, 153–172. Frankfurt am Main: Peter Lang.

Sudo, Yasutada. 2012. On the semantics of phi features on pronouns. Doctoral dissertation, MIT, Cambridge, MA.

Thornton, Anna M. 2001. Some reflections on gender and inflectional class assignment in Italian. In *Naturally! Linguistic studies in honour of Wolfgang Ulrich Dressler presented on the occasion of his 60th birthday*, ed. by Chris Schaner-Wolles, John Rennison, and Friedrich Neubarth, 479–487. Turin: Rosenberg & Sellier.

Thornton, Anna M. 2003a. La representazione dell'informazione morfologica nelle entrate lessicali. In *Scritti di morfologia in onore di Sergio Scalise in occasione del suo sesantesimo compleanno*, ed. by Antonietta Bisetto, Claudia Jacobini, and Anna M. Thornton, 203–221. Rome: Caissa Italia editore.

Thornton, Anna M. 2003b. L'assegnazione del genere in italiano. In *Actas del XXIII Congreso Internacional de Lingüística y Filología Románica*, ed. by Fernando Sánchez Miret, 467–481. Tübingen: Niemeyer.

Wechsler, Stephen, and Larisa Zlatić. 2000. A theory of agreement and its application to Serbo-Croatian. *Language* 76:799–832.

Wechsler, Stephen, and Larisa Zlatić. 2003. *The many faces of agreement.* Stanford, CA: CSLI Publications.

Welch, Nicolas. 2014. A tripartite agreement: Classificatory verbs, animacy and inflection in Tłı̨chǫ Yatıì. Paper presented at the annual meeting of the Linguistic Society of America, Minneapolis, 5 January 2014.

Williams, Edwin. 2003. *Representation Theory*. Cambridge, MA: MIT Press.

Wiltschko, Martina. 2009. What's in a determiner and how did it get there? In *Determiners: Universals and variation*, ed. by Jila Ghomeshi, Ileana Paul, and Martina Wiltschko, 25–66. Amsterdam: John Benjamins.

Wiltschko, Martina, and Elizabeth Ritter. 2015. Animating the narrow syntax. *The Linguistic Review* 32: 869–908.

Wiltschko, Martina, and Olga Steriopolo. 2007. Parameters of variation in the syntax of diminutives. In *Proceedings of the 2007 Annual Conference of the Canadian Linguistic Association*, ed. by Milica Radišić. http://cla-acl.ca/actes-2007-proceedings.

Yatsushiro, Kazuko, and Uli Sauerland. 2006. [feminine] in a high position. *Snippets* 13:11–12.

Linguistics and Languages
McMaster University
TSH-509
1280 Main Street West
Hamilton, Ontario L8S 4M2
Canada

ivona@alum.mit.edu