Universitat d'Alacant
Universidad de Alicante

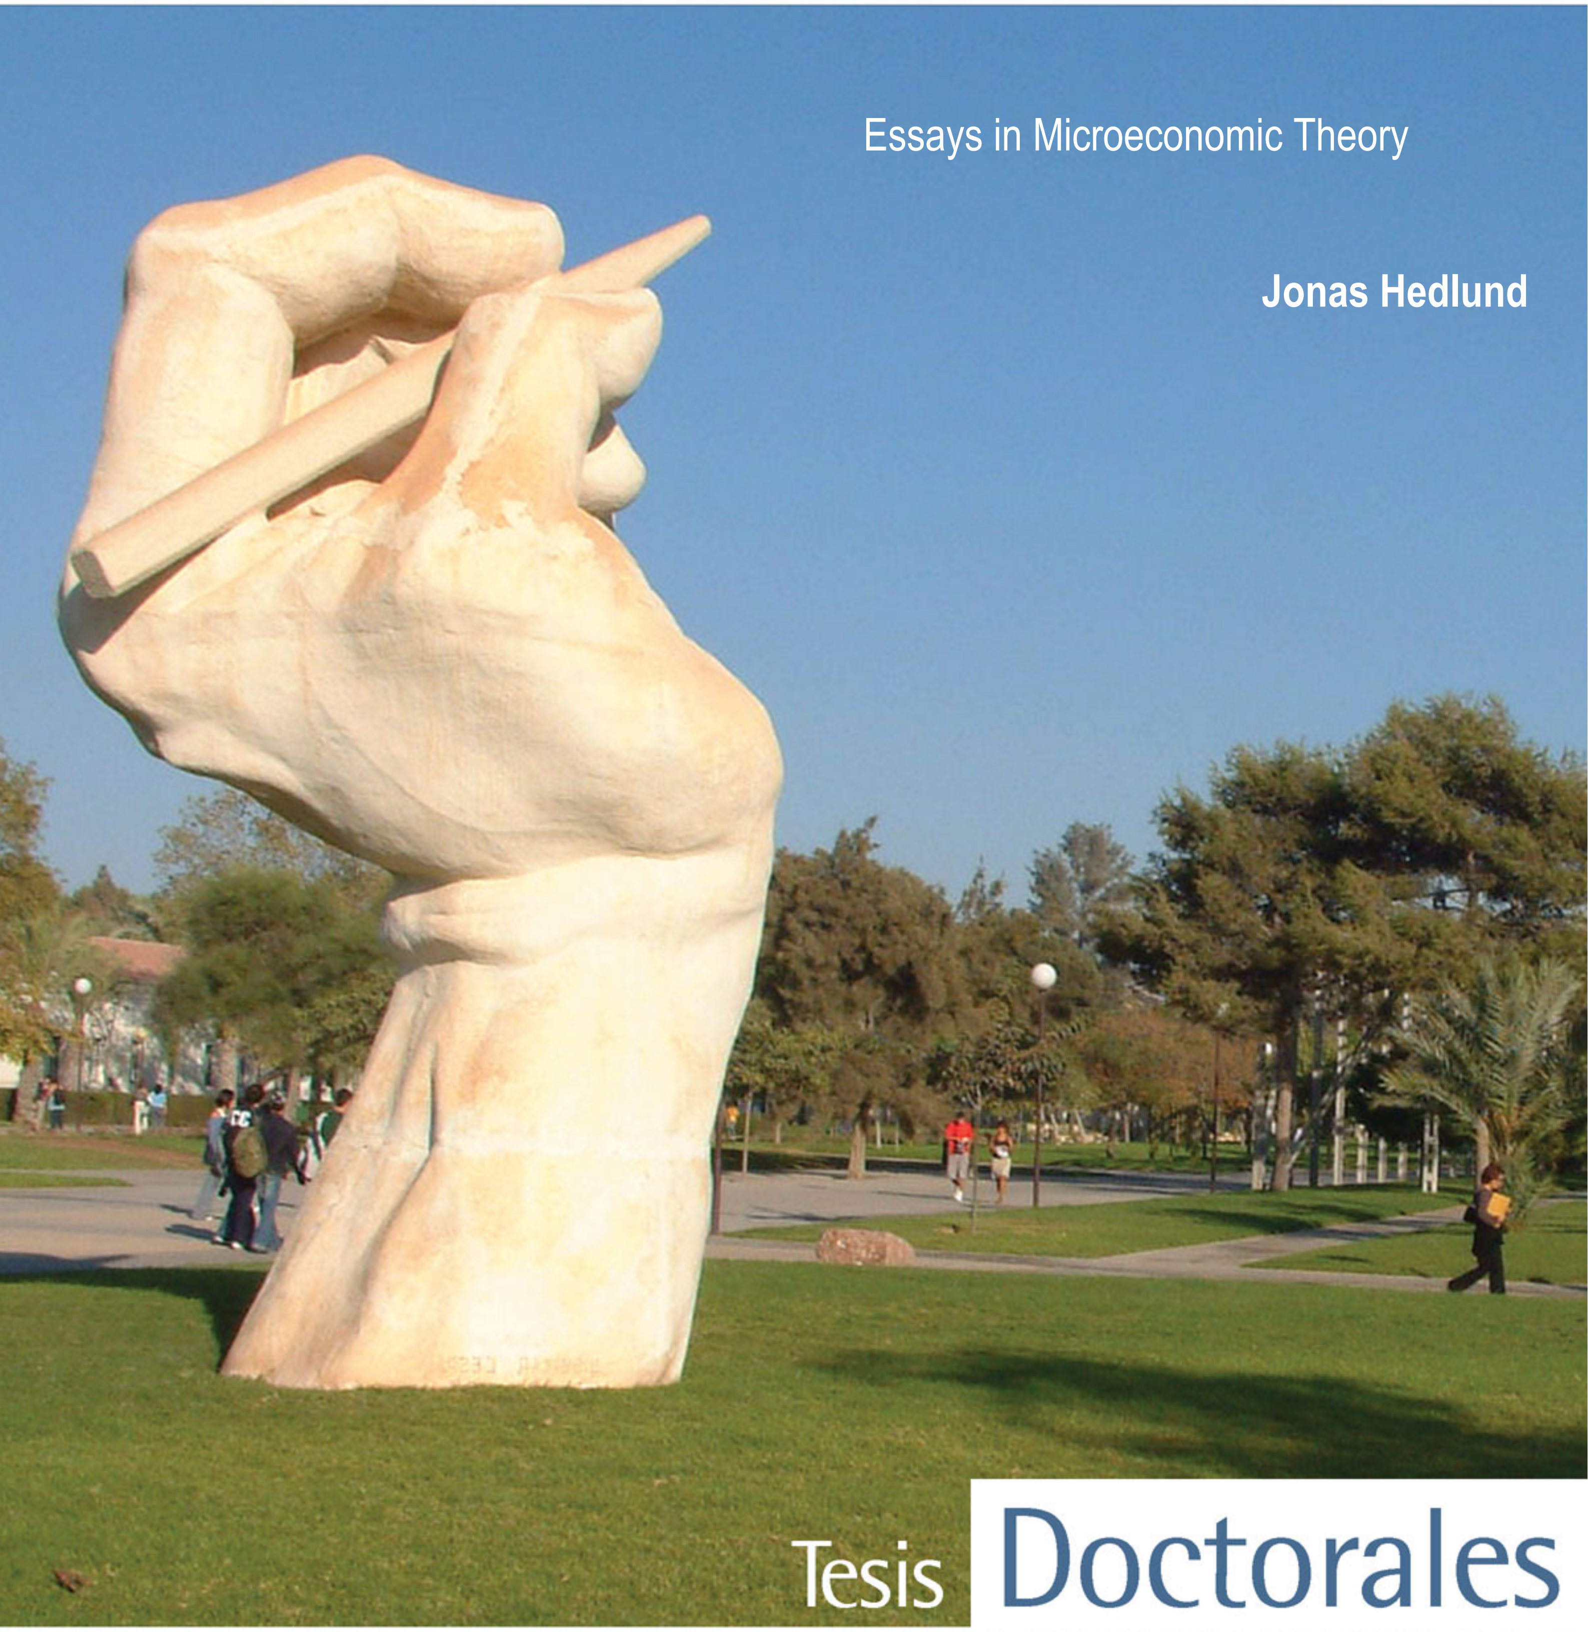

Essays in Microeconomic Theory

Jonas Hedlund

Tesis Doctorales
www.eltallerdigital.com
UNIVERSIDAD de ALICANTE

Essays in Microeconomic Theory

Jonas Hedlund

Supervisor: Luis Ubeda

Quantitative Economics Doctorate
Departamento de Fundamentos del Análisis Económico
Universidad de Alicante

April 2011

To my Family

Agradecimientos

En primer lugar quiero agradecer a mi esposa Yatziry por el ánimo y apoyo constante, por su paciencia, comprensión y flexibilidad, durante todo el trancurso de este proyecto. Sin su gran disposición, no hubiera sido posible. Agradezco a mi familia - Alf, Mona, Niklas, Lina, Johan, P-O y mis abuelos Åke y Barbro - por siempre estar al pendiente y brindarme apoyo de todo tipo. Un gracias especial a mi padre, Alf, por haber inspirado en mi el amor al conocimiento.

Agradezco a mi supervisor Luis Úbeda por todos los consejos, y todo el apoyo y ayuda, los cuales han sido cruciales en cada paso de la elaboración de la tesis, y que a menudo me han dado claridad donde antes había confusión. Un gracias especial a Carlos Oyarzun por apoyo, consejos y por comentarios que han ayudado a mejorar los manuscritos de manera significativa. Agradezco también a Carlos por trabajo conjunto en el primer capítulo de la tesis. Gracias a Adam Sanjurjo y a Íñigo Iturbe por comentarios muy útiles sobre la tesis.

Estoy también muy agradecido con todo el personal académico y no académico del departamento de Fundamentos del Análisis Económico de la Universidad de Alicante. Quiero además agradecer a los profesores Paul Nystedt y Victor Carreón, quienes con su excelencia docente motivaron mi interés por la economía. Agradezco a Maite Guijarro por ayudar a darme la posibilidad de estudiar la maestría en economía en el CIDE. Finalmente le doy las gracias a todos mis compañeros del doctorado, pero en particular a Gergely y a Gustavo, tanto por pláticas sobre temas de investigación, como por los buenos momentos que hemos compartido a través de estos años .

Contents

Introducción

La corriente principal de la teoría económica se basa en el supuesto de que los agentes económicos son racionales y maximizan la utilidad. Por ejemplo, diferentes situaciones económicas se modelan a menudo suponiendo que los agentes poseen una funcion de utilidad sobre el conjunto de alternativas factibles y que escogen la alternativa que maximiza el valor esperado de esta function. Luego, un equilibrio es típicamente definido como una situación en la que todos los agentes se comportan de esta manera al mismo tiempo y nadie tiene incentivos para modificar su decision. El equilibrio constituye la predicción del modelo.

Sin embargo, un cuerpo creciente de evidencia muestra que este modelo de comportamiento no siempre es una buena aproximación de como los agentes hacen sus decisiones. En vez de esto, resulta que muchas veces hacen sus decisiones de una manera más heuristica, por ejemplo recurriendo a simples "reglas de oro" de decisión, consistentes con el concepto de la racionalidad acotada (véase por ejempo Conlisk 1996, Gigerenzer, Todd and the ABC Research group 1999). Esto implica que los equilibrios de los modelos de maximización de utilidad no necesariamente son buenas predicciones del comportamiento de los agentes económicos.

Existe una literatura amplia que, tomando nota de este problema, estudia el impacto a las predicciones de los modelos económicos si los agentes en vez de maximizar la utilidad hacen sus decisiones utilizando reglas heurísticas de decisión. Una corriente de esta literatura que se ha ido incorporando en la teoría económica a lo largo de los últimos veinte años se conoce como la *teoría evolutiva de juegos.*[1] En la teoría evolutiva de juegos es comun suponer que los agentes ajustan su comportamiento aplicando una regla de decision de racionalidad acotada. Luego se analiza como la sociedad evoluciona a través del tiempo bajo este supuesto. Una característica atractiva de la teoría evolutiva de juegos es su dimensión dinámica. En varios modelos económicos estándares el equilibrio es una noción estática. Se considera que la ausencia de incentivos para modificar las decisiones implica que no hay fuerzas que mueven la economía en una u otra dirección. Sin embargo, no se hace mucha mención del proceso por el cual la economía llega a ese equilibrio[2]. Al contrario, una característica típica de la teoría evolutiva de juegos es que modela este proceso explícitamente. De esta manera, la teoría evolutiva de juegos permite analizar si la economía realmente converge hacía el equilibrio y cual sería la trayectoria por la que llega allí.

Esta tesis contribuye a la teoría evolutiva de juegos explorando el impacto de *reglas de decision imitiativas* sobre las predicciones de ciertos modelos económicos. La idea fundamental es que los agentes observan lo que otros agentes hacen y copian

[1] Véase por ejemplo Sandholm 2010; Weibull 1995. Algunas contribuciones cruciales de esta literatura son Young 1993; Kandori, Mailath y Rob 1993; Young 1998.

[2] Fudenberg y Levine 1998.

sus decisiones en funcion de cuan exitosos parecen ser. Por tanto, los agentes utilizan la información contenida en el comportamiento de los demás para hacer decisiones. El comportamiento imitativo es un ejemplo de toma de decision bajo racionalidad acotada. Simplemente requiere observar lo que hacen los demás y a lo mejor llevar a cabo alguna operación matemática elemental. Existe evidencia amplia que este tipo de comportamiento es una parte prominente de la toma de decisiones de los seres humanos. Por ejemplo, se considera fundamental en ciertas ramas de la psicología (véase por ejemplo Bandura 1977). También hay una literatura experimental que estudia el comportamiento imitativo en situaciones de toma de decisiones económicas y que confirma su importancia (por ejemplo Huck, Normann and Oechssler 1999 y 2000; Apesteguia, Huck and Oechssler 2007 y 2009; Pingle y Day 1996; Offerman and Schotter 2009). Dado que los agentes económicos a menudo hacen sus decisiones de esta manera, es importante estudiar como el comportamiento imitativo afecta las predicciones de los modelos económicos.

El método empleado para realizar el estudio es el siguiente: Se considera un conjunto de modelos económicos emblemáticos que constituyen puntos de referencia en la teoría económica. Para cada uno de ellos se especifica la forma exacta en la que los agentes observan el comportamiento de los demás, y la manera en la que utilizan esta información para hacer decisiones. Finalmente se estudia el proceso dinámico que surge como consequencia de este comportamiento. En particular, en cuanto a metodología el análisis depende en gran medida de las herramientas matemáticas desarrolladas por Freidlin and Wentzell (1986) para el análisis de procesos estocásticos perturbados, e introducidas a la economía por Young (1993) y Kandori, Mailath y Rob (1993).

Existe una literatura teórica sustancial sobre dinámicas de imitación en la economía. Por ejemplo, Vega-Redondo (1997) inició una literatura sobre el impacto de dinámicas de imitación en oligopolios. Schlag (1998) y Oyarzun y Ruf (2009) caracterizaron reglas de decisión imitativas que siempre llevan a una población a escoger la alternativa que maximiza la utilidad en el largo plazo. Ellison y Fudenberg (1992) estudiaron la adopción de tecnología a través de la imitación. Eshel, Samuelson y Shaked (1998) mostraron que el comportamiento imitativo e interacción local en conjunto pueden llevar a resultados cooperativos en situaciones en las que los tomadores de decisiones racionales nunca cooperarían. En particular, esta tesis hace tres contribuciones a la literatura sobre dinámicas de imitación. El primer capítulo estudia la imitación en poblaciones heterogeneas[3]. La dificultad que surge en una población heterogenea es que lo que funciona bien para un agente puede no ser bueno para otro. Esto implica que el desempeño de reglas imitativas de decision es ambiguo. Se muestra que en una población heterogenea la imitación tiende a llevar a la sociedad hacia resultados que son buenos para la fracción que constituye la mayoría, pero que

[3]Este capítulo se ha elaborado en colaboración con Carlos Oyarzun.

perjudica a los grupos minoritarios. La razón es que para la minoría hay una fuerte tendencia de encontrarse con agentes de otro tipo, quienes escogen lo que es bueno para ellos pero malo para la minoría. En el segundo capítulo se analiza la imitación en oligopolios, en situaciones en las que hay varios mercados y las empresas a veces observan empresas en otros mercados. Se muestra que en estos casos las predicciones del modelo son menos competitivos que en el caso de un único mercado (por ejemplo Vega-Redondo 1997). La intuición es que un mercado aislado en el que prevalece la cooperación tiende a ser imitado por empresas en otros mercados. En el tercer capítulo se demuestra que el supuesto de interacción local en Eshel et. al. (1998) que garantiza que la cooperación pueda persistir en el largo plazo, en realidad es menos restrictivo de lo que estos autores encontraron. Mientras que la imitación siga siendo local, la cooperación puede prevalecer aún cuando las decisiones de cada agente afecten a casi toda la población.

Finalmente, el cuarto capítulo de la tesis hace una contribución a la literatura de transmisión estratégica de información. Muchas situaciones económicas se caracterizan por la información privada. A veces esto puede causar ineficiencias. Por ejemplo, Akerlof (1970) demostró con un ejemplo del mercado de coches usados que el mercado entero puede desaparecer cuando el vendedor tiene más información que el comprador. Sin embargo, Milgrom (1981) mostró que si es posible la comunicación, puede haber una solución a este problema. Por ejemplo, si el vendedor puede proveer un reporte certificable con respecto a diferentes aspectos del coche, las asimetrías de información pueden ser eliminadas. En particular, Milgrom (1981) mostró que si el vendedor puede comunicar toda su información sin costes, siempre lo va a hacer. No obstante, reportar información privada a menudo es costoso. La pregunta es hasta dónde la comunicación puede resolver el problema de la información asimétrica en presencia de costes. En el cuarto capítulo de esta tésis se demuestra que la comunicación siempre puede eliminar todas las asimetrías de información, aún cuando los costes de la comunicación son arbitrariamente altos. La intuición es que siempre es posible para el agente informado identificarse completamente con un reporte tal que otros agentes no tratarían de hacerse pasar por él, o porque no disponen de la misma información o porque los costes the producirla son demasiado altos.

A continuación se describen cada uno de los capítulos de la tesis con más detalle.

Capítulo 1: Aprendizaje por comparación social en una población heterogénea

La idea subyacente de la la imitación como regla de decisión es que los individuos pueden utilizar la información contenida en las experiencias de los demás para hacer buenas decisiones. Si los individuos son similares, esto puede llevar a una población a escoger la mejor acción en el largo plazo (Ellison y Fudenberg 1995; Schlag 1998). Sin embargo, en una población heterogénea surge un problema, ya que lo que funciona bien para un individuo puede no ser una buena alternativa para otro.

En este capítulo se incorpora una noción de *comparación social* (véase por ejemplo Festinger 1954 o Suls, Martin y Wheeler 2000) en un modelo de aprendizaje de racionalidad acotada. La teoría de la comparación social postula que cuando los individuos no pueden evaluar sus habilidades por información objetiva, pueden hacer una evaluación relativa con respecto a otros individuos. Suponemos que los individuous utilizan la comparación social para circunvenir el problema de la heterogeneidad y extraer información relevante de las experiencias de los demás. Analizamos una población de individuos que repetidamente escogen una acción del mismo conjunto. Como consequencia de su decisión, cada individuo obtiene un pago de acuerdo a alguna distribución de pagos. Además, estas distribuciones pueden ser diferentes para diferentes individuos. El flujo de información que consideramos es parecido al modelo canónico en el que los tomadores de decisiones pueden observar la experiencia de los demás, es decir, las acciones que escogen y los pagos que obtienen (por ejemplo Fudenberg y Ellison, 1993, 1995; Vega-Redondo 1997; Schlag 1998, 1999). En particular, cada individuo observa a otro individuo seleccionado aleatoriamente de la población. Por lo tanto, en cada periodo el individuo observa la acción que escogió, el pago que obtuvo y la acción y el pago de algún otro individuo. Incorporamos la noción de comparación suponiendo que cada individuo además percibe una señal aleatoria que informa sobre la diferencia en el pago esperado del individuo y el individuo observado con respecto a la acción del individuo observado. Suponemos que la señal es no sesgada, es decir, los individuos no se equivocan sistemáticamente en su evaluación de la señal.

Los individuos utilizan toda esta información para determinar su comportamiento en el siguiente periodo. Llamamos *regla de decision* a la función que asigna una probabilidad para escoger cada acción dados las acciones observadas, los pagos obtenidos y la señal de similitud percibida. Nos enfocamos en unas reglas de decisión que satisfacen ciertas propiedades en cuanto a su desempeño, medido por el pago esperado de la acción que escoge. En particular, nuestro análisis empieza con la caracterización de una clase de reglas de decisión tal que el individuo siempre escoge la mejor de las acciones de su muestra con mayor probabilidad. Llamamos a a esta clase de reglas de decisión *monotonas en el pago.* Nuestro primer resultado (Proposición 1) muestra que para reglas de decisión monótonas en el pago, la probabilidad actualizada de escoger la acción observada es una transformación afín de la diferencia entre la suma del pago obtenido por el individuo observado y la señal de similitud y el pago obtenido con la acción escogida. Las reglas de decisión monótonas en el pago permiten a los individuos hacer decisiones tales que la heterogeneidad *per se* no les engaña.

Nuestra preocupación principal son las dinámicas de las decisiones de los individuos cuando todos utilizan una regla de decisión monótona en el pago. Consideramos una población con dos tipos de individuos, A y B, que pueden escoger entre dos acciones distintas, a y b. Las distribuciones de pago asociada a cada una de las acciones son las mismas para individuos del mismo tipo y el pago esperado de la acción a (b) es mayor que el pago esperado de la acción b (a) para todos los individuos de tipo A (B). Nuestros resultados principales (Proposición 2 y 3) demuestran que la decisiones de los individuos convergen a un punto estable determinado por la fracción de cada tipo en la población y la diferencia en ganancia, en términos esperados, de escoger la acción óptima para los diferentes tipos de individuos. En particular, si las fracciones de los dos tipos de la población son suficientemente diferentes, entonces toda la población converge a la acción óptima del tipo mayoritario, y por lo tanto, a la acción subóptima de la minoría. La intuición es que tanto la mayoría como la minoría se mueven hacia sus acciones óptimas respectivas. Sin embargo, el movimiento del tipo mayoritario es más fuerte y este movimiento ejerce una influencia poderosa sobre la minoría. Por consiguiente, la minoría puede acabar escogiendo la acción equivocada, a pesar de utilizar una regla de decisión que le permite escoger la acción óptima con mayor probabilidad.

La acción a la que la población converge también es determinada por las ganancias en términos esperados de escoger la acción óptima. Si los tamaños de las poblaciones de cada tipo son parecidos, y estas ganancias son suficientemente diferentes, entonces toda la población converge a la acción óptima del tipo que se beneficia más cambiándose a la acción óptima. Intuitivamente, la aplicación de la regla monótona en pagos hace que las acciones óptimas de cada tipo se propaga, y la acción óptima que reporta la mayor ventaja sobre la acción subóptima se propaga a mayor velocidad.

El capítulo contribuye a la literatura sobre dinámicas de imitación. Los artículos seminales de esta literatura son Ellison and Fudenberg (1993, 1995), Vega-Redondo (1997), and Schlag (1998)[4]. La mayoría de estos trabajos no consideran la heterogeneidad y el hecho de que los individuos pueden ser concientes de ella y ajustar sus reacciones de acuerdo a ello. Una excepción notable es Ellison y Fudenberg (1993). En su trabajo los individuos están restringidos a observar solamente los vecinos, quienes en la mayoría de los casos comparten la misma acción óptima. De hecho, cuanto más pequeña sea la "ventana" de vecinos que cada individuo observa, más cerca estará el sistema a la convergencia a la alternativa óptima para toda la población. En contraste, nuestro análisis no supone que los individuos tienden a observar más a menudo individuos que son parecidos a ellos. En nuestro modelo, esto permite que en particular las minorías terminan escogiendo una acción subóptima.

[4] Para una revisión completa de esta literatura véase Alós-Ferrer y Schlag 2009.

Capítulo 2: Imitación en oligopolios Cournot con mercados múltiples

En los oligopolios de Cournot, el supuesto de que los agentes toman sus decisiones maximizando una funcion objetivo se traduce en un supuesto de maximización de beneficios. Sin embargo, en años recientes una literatura ha emergido que estudia decisiones basadas en la imitación más bien que la maximización de los beneficios. Esto tiene sentido, ya que los oligopolios son situaciones complejas en las que es probable que las empresas a veces recurren a procesos de decisión mas sencillas que la maximización de los beneficios. Además hay varios artículos experimentales que confirman la importancia de la imitación en los oligopolios[5].

La imitación en oligopolios Cournot a veces lleva a resultados muy competitivos. Por ejemplo, si las empresas en un mercado en cada periodo escoge la cantidad que generó los beneficios más altos del mercado en el periodo anterior, entonces la predicción a largo plazo corresponde a la de la competencia perfecta (Vega-Redondo 1997). Sin embargo, esto no se cumple necesariamente si hay más de un mercado y las empresas a veces imitan a través de los mercados (Apesteguia, Huck and Oechssler, 2007). Resulta que la información sobre el desempeño de las empresas en otros mercados es una variable importante en este caso. Sin embargo, no hay muchos estudios sobre el impacto de la información a las predicciones del modelo. Este capítulo contribuye al estudio de la imitación en los oligopolios Cournot, a través del análisis riguroso de como la información sobre el desempeño de las empresas en otros mercados afecta a los mercados en el largo plazo.

Supondré que hay un conjunto de mercados idénticos y separados. En cada periodo las empresas obervan a las empresas del mismo mercado y una muestra de empresas de los otros mercados. Luego escogen una cantidad a través de una regla de decisión. Se consideran dos reglas de decisión distintas. Si una empresa utiliza *imita el mejor máximo* (IBM, por sus siglas en inglés) escoge la cantidad que generó el beneficio más alto en la muestra. Sin embargo, cuando hay varios mercados, una cantidad puede generar beneficios altos en un mercado y bajos en otros. En este caso no está claro que sería atractivo imitar esa cantidad. Si una empresas utiliza *imita el mejor promedio* (IBA, por sus siglas en inglés) toma esto en cuenta al calcular los beneficios promedio de cada cantidad de su muestra y luego imitar la que generó el beneficio promedio más alto. También se permite que las empresas de vez en cuando experimentan y escogen una cantidad al azar. Utilizo las técnicas desarrolladas por Young (1993) para caracterizar el soporte de la distribución estacionaria de la resultante cadena de Markov perturbada. Este soporte se conoce como los estados *estocásticamente estables* y constituyen la predicción a largo plazo de las dinámicas.

Se consideran tres escenarios informacionales distintos. En el primero, como

[5] Por ejemplo, Huck, Normann y Oechssler 1999, 2000; Offerman, Potters y Sonnemans 2002; Apesteguia, Huck and Oechssler 2007, 2009.

punto de referencia, las empresas observan todas las empresas en todos los mercados. En este caso, cuando todas las empresas utilizan IBM todas las cantidades entre la de Cournot y la competencia perfecta aparecen en el largo plazo. La intuicion es que las cantidades cercanas a la de competencia perfecta tienden a tener el mejor desempeño en cada mercado. Al mismo tiempo, los mercados que producen cantidades cercanas a la de Cournot tienden a tener un mejor desempeño que otros mercados. Esto crea tendencias que van en direcciones contrarias, los cuales se equilibrian en el largo plazo.

Cuando las empresas utilizan IBA, existe un único estado estocásticamente estable, en el que todas las empreasa producen una cantidad estrictamente entre la cantidad de Cournot y la de competencia perfecta. Esta cantidad decrece en el número de mercados y crece en el número de empresas por mercado. El resultado aparece ya que el número de empresas por mercado y el número de mercados determinan ponderadores que afectan el cómputo de los beneficios promedios. Por ejemplo, un número más grande de mercados aumenta la importancia relativa de las observaciones a través de los mercados. Se muestra que el estado estocástimente estable corresponde al único equilibrio simétrico de Nash de un juego en el que a las empresas les importan tanto sus ganancias en términos absolutos como en comparación con otras empresas.

En el segundo escenario informacional, las empresas observan empreas en mercados en los que la cantidad agregada es suficientemente parecida a la del propio mercado. La idea es que la imitación intuitivamente tiene más sentido si las empresas observadas están en una situación similar a la propia. Si a las empresas no les importan mucho las diferencias en el agregado, el modelo se comporta como en el caso de información global. Si son muy sensibles con respecto a diferencias en el mercado, obtenemos competencia perfecta (como en Vega-Redondo 1997). En casos intermedios el resultado es más competitivo cuando las empresas son sensibles a diferencias en el agregado. Intuitivamente, la cautela en cuanto a imitar empresas en otros mercados tiende a aislar a los mercados, y esto les lleva a evolucionar independientemente.

En el tercer escenario informacional los mercados están ordenados en un círculo y las empresas observan el propio mercado y las empresas de algunos mercados vecinos. Esto refleja una situación en la que la localización geográfica de las empresas impide que observen todos los demás mercados. Los resultados del caso de información global son robustos a este escenario. Esto implica que si las empresas utilizan IBA, el resultado es menos competitivo cuando las empresas observan una cantidad mayor de mercados.

Una conclusión que es válida a través de los diferentes escenarios es que más información conduce a resultados menos competitivos. Este debe ser contrastado con la conclusión de Huck et. al. (1999, 2000), que más información sobre las empresas del mismo mercado lleva a resultados más competitivos. Cuando hay varios mercados, más información sobre las empresas en otros mercados puede conducir a resultados menos competitivos.

Este capítulo es relacionado a Apesteguia et. al. (2007), quienes fueron los primeros en descubrir que la información de empresas en otros mercados afecta el resultado. Sin embargo, el modelo considerado aquí se distingue de Apesteguia et. al. (2007) de varias maneras. Mientras que Apesteguia et. al. (2007) consideran un modelo lineal sin costes de producción, aquí solo suponemos que la demanda es decreciente y cóncava y que los costes son crecientes y convexos. En Apesteguia et. al. (2007) la capacidad de cada empresa es limitada a la cantidad de competencia perfecta y cada empresa es asignada aleatoriamente a un nuevo mercado en cada periodo. Aquí, no suponemos que la capacidad es limitada y los mercados no se reasignan de un periodo a otro. Las conclusiones que obtenemos son claramente diferentes. Por ejemplo, en el caso de información global, aquí se predice que todas las cantidades entre la de Cournot y la de competencia perfecta pueden aparecer en el largo plazo. Apesteguia et. al. (2007) obtiene competencia perfecta para el caso correspondiente. En el caso de IBA, aquí se predice una cantidad única estrictamente entre la de Cournot y la de competencia perfecta, mientras que Apesteguia et. al. (2007) obtienen la de Cournot. Por lo tanto, este capítulo puede verse como un estudio ampliado de las dinámicas descubiertas por primera vez en Apesteguia et. al. (2007). De esta manera se muestra que cuando se relajan ciertos supuestos, se obtienen diferentes conclusiones.

Este capítulo también es relacionado a modelos en los que las empresas se acuerdan del pasado (Alós-Ferrer 2004; Bergin y Berghardt 2009). Esta relación tiene que ver con el hecho de que el modelo con memoria puede considerarse como un modelo de varios mercados, en el que los mercados adicionales existen en la memoria de las empresas.

Capítulo 3: Provisón altruista de bienes públicos e interacción local

Un supuesto estándar en la teoría económica es que los individuos actúan en su mejor interés. Sin embargo, muchas veces se observan a las personas actuando más allá de sus propios intereseses, algo que es conocido como *altruismo*. La teoría económica ha generado un número de modelos para reconciliar los actos altruistas con la racionalidad típica de muchos modelos económicos. Por ejemplo, un enfoque ha sido modelar el altruismo modificando la función de utilidad (como en Becker 1974 y 1981), o suponer que la interacción se repite infinitamente (véase por ejemplo Fudenberg y Maskin 1986).

Recientemente, una explicación alternativa ha sido propuesta por Eshel, Samuelson y Shaked (1998). Si en una población circular los individuos altruistamente comparten un bien público con los vecinos más cercanos y además hacen sus decisiones imitando a esos mismos vecinos, entonces el comportamiento altruista puede persistir. Este resultado es robusto aun en presencia de experimentaciones o "mutaciones" de parte de los individuos. La intuición es los altruistas se pueden juntar en grupos, excluyendo a los egoistas. De esta manera obtienen mayores pagos y tienden a ser imitados.

El supuesto de que la imitación es local es razonable. La imitación require conocimiento de lo que los demás hacen y como les va. Este tipo de información la tenemos generalmente sole de un pequeño subconjunto de los miembros de la sociedad. Al otro lado, muchas externalidades son mucho menos locales. Por ejemplo, tirar basura en la calle afecta a personas que ni siquiera conocemos. En particular es probable que afecte a más individuos de las que imitamos.

Por esta razón, en este capítulo, analizo un modelo en el que se permito que el bien público sea menos local que la imitación. En cada periodo cada individuo decide si proveer un bien público (y ser altruista), o no hacerlo (y ser egoista). El bien público es compartido con un número arbitrario de vecinos. Para hacer una decisión, el individuo observa la acción y el pago de los vecinos más cercanos. Luego escogen la acción que generó el pago más alto en promedio. Además se permite que los individuos de vez en cuando experimenten y escojan una cantidad al azar. Se muestra que en ausencia de experimentaciones, el altruismo puede prevaleceer y coexistir con el egoismo siempre y cuando el bien públio no sea global. En la presencia de experimentación, el altruismo puede persistir, pero solo si la población es suficientemente grande. Por lo tanto, la conclusión es que el resultado de Eshel et. al. (1998) es relativamente robusto a situaciones en las que la imitación es más local que el bien público.

La intuición es que el hecho que el bien público es local permite que los altruistas se junten en grupos, excluyendo así a los egoistas. De esta manera los altruistas obtienen un mayor pago y por tanto tienden a ser imitados. Al hacerse menos local el bien público, los altruistas tienen que juntarse en grupos más grandes para excluir a los egoistas en suficiente grado. Esto implica que al ser casi global el bien publico, al máximo un grupo altruista cabe en la población. Resulta que una constelación así es muy sensible a la experimentación y es suficiente que un solo altruista cambie a egoismo para que la población descienda a un estado de egoismo. Por esta razón cuando el bien público se comparte con más vecinos, se require una población grande para que el altruismo pueda persistir en presencia de experimentaciones.

La investigación se relaciona con Jun y Sethi (2007), Matros (2008) y Mengel (2009). Jun y Sethi (2007) y Matros (2008) consideran un modelo en el que el bien público es compartido con exactamente el mismo (potencialmente grande) conjunto de individuos que cada individuo imita. Por lo tanto, en estos artículos la estructura de interacción es menos estilizada que en Eshel et. al. (1998). No obstante, cuando el bien pública es casi global, no es probable que se cumpla el supuesto de que el conjunto de individuous que comparte el bien público coincide con el conjunto que cada individuo imita. Mengel (2009) supone que la imitación es menos local que el bien público. Encuentra que en este caso el altruismo no persiste en el largo plazo. Sin embargo, como ya se mencionó, es probable que en varios casos, sucede lo contrario. Mientras que los efectos externos de nuestras acciones muchas veces afectan a un número grande de individuos, a menudo imitamos a un conjunto más limitado.

Capítulo 4: Comunicación perfecta con costes de comunicación arbitrarios

La comunicación estratégica es un aspecto importante en muchas situaciones económicas en las que alguna parte tiene información privada. Los llamados *juegos de persuasión*, introducidos por Paul Milgrom en 1981 en un artículo influyente, se enfoca en la comunicación estratégica en términos de información verificable. En su modelo, un *emisor* pretende influir en el comportamiento de un *receptor*, comunicando información por medio de un reporte. Milgrom (1981) mostró que si la comunicación no tiene costes y las preferencias sastisfacen una condición de monotonicidad, entonces el emisor revela toda su información privada al receptor. Es decir, la comunicación es perfecta en la presencia tanto de preferencias incongruentes y comportamiento estratégico.[6]

Sin embargo, mientras Milgrom (1981) supuso que la comunicación no tiene costes, reportar información verificable muchas veces es costoso. A menudo requiere explicaciones cuidadosas y detalladas basadas en hechos. Por ejemplo, un emprendedor que elabora un plan de negocio para convencer a un inversor gasta tiempo y esfuerzo explicando los detalles técnicos del producto, los costes de producción y demás información relevante. Además, puede implicar certificación costosa de instituciones acreditadas, como auditoría por parte de contadores externos, o un patente que acreditar la originalidad del producto. Al mismo tiempo, la literatura existente indica que el resultado de Milgrom (1981) es sensible al supuesto de que la comunicación no tiene costes (véase Jovanovic 1982; Verecchia 1983 y Cheong y Kim 2004).

En este trabajo se llega a una conclusión diferente. Estudio un modelo basado en el juego de persuasión de Milgrom (1981), en el que el coste de reportar información privada crece continuamente en la precisión del reporte. La conclusión es que el resultado de Milgrom (1981) es relativamente robusto a los costes de comunicación. Más específicamente, mientras que Milgrom (1981) mostró que si la comunicación no tiene costes, entonces cada equilibrio es separador, aquí se muestra que un equilibrio separador siempre existe, sin importar los costes. La intuición es que los costes de

[6] Diferentes generalizaciones del resultado de Milgrom 1981 se pueden encontrar en Milgrom y Roberts (1986), Seidman y Winter (1991) y Mathis (2008).

comunicación introducen señalamiento costoso al juego de persuasión. El reporte revela información y al mismo tiempo los costes funciona como un dispositivo de señalamiento. Cuando es demasiado caro reportar toda la información, un emisor de tipo alto puede impedir que los tipos más bajos copien su reporte a través de los costes. Resulta que una combinación de revelación de información y señalamiento costoso siempre puede lograr la separación completa

Se muestra que pueden haber varios equilibrios separadores, pero todos ellos son equivalentes en términos de los pagos. Cuando los costes son bajos, todos los equilibrios son separadores. La unicidad del equilibrio separador aparece ya que con costes bajos, los tipos emisores altos siempre pueden desviarse de cualquier equilibrio agrupador revelando su tipo verdadero. Cuando los costes son altos, aparece un equilibrio. En este equilibrio, los reportes no contienen ninguna información. La intuición es que con costes altos, el escepticismo de parte del receptor puede hacer que sea demasiado caro para cualquier emisor revelar su tipo verdadero.

La razón por la que Jovanovic (1982), Verecchia (1983) y Cheong y Kim (2004) llegan a una conclusión distinto, es que en sus modelos el emisor tiene que reportar toda su información, o ninguna. Así, siempre hay algunos tipos de emisores que no revelan su información, y si el coste es demasiado alto, el emisor nunca revela ninguna información en absoluto. En este trabajo, el emisor tiene más discreción en este respecto. Decide continuamente cuanta información reportar. Por lo tanto, si los costes son altos, puede reportar una cantidad pequeña de información sin incurrir demasiados costes.

Finalmente, se introduce una extensión al juego de persuasión en el receptor tiene que haer un esfuerzo costoso para acceder la información contenida en el reporte. La idea es tomar en cuenta el hecho de que muchas veces es costoso leer y entender un reporte. Cuando el receptor decide activamente si lee un reporte, puede condicionar su decisión a la primera impresión de él. Una cuestión relevante es como se relaciona esta primera impresión con el contenido del reporte y en que grado el emisor puede manipularlo. Aquí supongo que el aspecto de un reporte se relaciona a la cantidad de información contenida en ella y que los aspectos pueden ser manipulados, incurriendo un coste. Se caracterizan dos clases de equilibrios separadores. En una de ellas, *equilibrios sin lectura*, el receptor nunca lee ningun reporte. Dependiendo de los costes, el emisor puede no manipular el aspecto del reporte, o invertir recursos para hacer que se parezca más preciso de lo que es. En la otra clase de equilibrios, *equilibrios de lectura*, el receptor lee todos los reportes. Como los dos equilibrios son separadores, el receptor prefiere el primero, en el que no hace ningún esfuerzo para leer los reportes.

Este capítulo se relaciona con Mathis (2008), quien considera comunicación en términos de información *parcialmente* verificable. La información parcialmente verificable se puede tratar como un modelo de comunicación costosa en el que ciertos reportes no tienen costes y otros tienen costes arbitrariamente altos. En otras palabras, los costes son discontinous. Mathis (2008) encuentra que si es arbitrariamente costoso revelar el tipo verdadero, no existe ningún equilibrio separador, lo cual contrasta con la conclusión de esta trabajo. La diferencia se debe a que en Mathis (2008) los costes son discontinuos. Por tanto, al tratar la verificabilidad parcial en términos de comunicación costosa, la continuidad de los costes tiene un impacto importante en el resultado.

Otro trabajo relacionado es Kartik (2009), quien considera mentiras costosas. En el modelo de Kartik (2009), el emisor puede reportar información falsa a un coste. Sin embargo, en contraste con el trabajo presente, el coste no se relaciona con la precisión del reporte. Kartik (2009) encuentra que la separación completa es imposible y caracteriza equilibrios en los cuales tipos bajos se separan y tipos altos se agrupan.

Introduction

This thesis contains the results of my doctoral studies in the Quantitative Economics Doctorate (QED) at the Universidad of Alicante. The title of the thesis is "*Essays in Microeconomic Theory*" and it consists of four different chapters. The focus of the thesis is on imitation dynamics in economic models and it also contains a chapter on strategic transmission of information.

The first chapter analyzes imitative behavior in a heterogenous population. The underlying idea is that an individual can use the information contained in others' choices and outcomes in order to make choices. If individuals are similar, this may lead a population to choose the optimal action in the long run (e.g., Ellison and Fudenberg, 1995; Schlag, 1998). However, in a heterogeneous population a problem arises, since what works well for one individual may not be a good choice for another. We incorporate a notion of *social comparison* (see for example Festinger 1954 or Suls, Martin and Wheeler 2000) which allows individuals to circumvent this problem and extract relevant information from the choices and outcomes of others. We analyze a population of individuals who repeatedly choose an action, thereby obtaining payoffs according to some unknown payoff distributions, which additionally may be different across individuals. In each period, each individual observes the action and payoff of some randomly selected individual, as well as a random signal that is informative about differences between the individual and the sampled individual. Individuals use this information to determine their behavior in the next period through a *decision rule*. We characterize *payoff monotone* decision rules, which allow the individual to on average choose the better among two sampled actions. This means that individuals can make choices such that heterogeneity *per se* does not mislead them.

The main concern is the dynamics of choices in a heterogenous population when each individual updates the probabilities of playing each action using a payoff monotone decision rule. We consider a population with two types of individuals, A and B, who choose between two actions, a and b. a (b) is optimal for type A (B) individuals. Our main results (Propositions 2 and 3) show that the choices of the population converge to a rest point determined by the fraction of each type in the population and the magnitude of the difference in payoffs between the optimal and suboptimal action for each type. If the fractions of the two types in the population are different enough, then the whole population converges to the majority type's optimal action. The intuition is that both the majority and the minority move toward their respective optimal actions. However, the motion of the majority type is stronger and since both types imitate each other this motion pulls the minority along. Hence, the minority may end up choosing the wrong action in spite of using a rule that makes the right choice, on average, among the actions observed. If the sizes of the population of each type are about the same, then the whole population converges toward the type that benefits the most from its optimal action. Intuitively, due to properties of the decision rule, the optimal action of the type that benefits the most from its optimal choice propagates faster, which leads the population toward this action in the long run.

The chapter contributes to the literature on imitation dynamics, with seminal papers Ellison and Fudenberg (1993, 1995), Vega-Redondo (1997), and Schlag (1998) (for a thorough survey, see Alos-Ferrer and Schlag 2009). Most of this work allows little role for heterogeneity and the fact that individuals may be aware of it and adjust accordingly. A notable exception is Ellison and Fudenberg (1993), in which individuals observe mostly neighbors who have the same optimal actions. The smaller the "window" of neighbors, the closer is the system to convergence to the optimal choice for the whole population. In contrast, our analysis does not assume that individuals tend to observe individuals that are more similar to them. This allows for the possibility that in particular minorities end up choosing a suboptimal action.

The second chapter of the thesis studies imitation in Cournot oligopolies when there are multiple markets. The study imitation of in oligopolies makes sense, since these are complex situations and it is likely that firms sometimes make decisions through processes that are cognitively simpler than profit maximization. There are also several experimental papers that confirm the importance of imitation in oligopoly games (Huck, Normann and Oechssler 1999, 2000, Offerman, Potters and Sonnemans 2002 and Apesteguia, Huck and Oechssler 2007, 2009).

In models of imitation dynamics in oligopolies, the information available about the performance of firms in other markets is an important variable (Apesteguia, Huck and Oechssler 2007). Still, there is not much literature on how different assumptions with respect to this information affect the outcome. Hence, the present paper provides a thorough analysis of how the information available about firms in other markets affect the long run behavior of the markets. I assume that there is a set of identical and separated markets and that in each period firms observe the firms in the own market and a sample of firms from the other markets. They then choose a quantity by using a decision rule. Either, they choose the quantity that generated the highest profit in the sample (*Imitate the Best Max,* IBM). Else, they compute the average profit of each observed quantity and imitates the quantity that generated the highest average profit (*Imitate the Best Average* (IBA)). Firms also sometimes experiment and choose a quantity randomly. I analyze the resulting dynamic process by using the techniques developed by Young (1993) to characterize the *stochastically stable states*, which constitutes the long run prediction of the dynamics.

Three different informational settings are considered. First, in the benchmark case, firms observe all firms in all markets. When firms use IBM all quantities between the Cournot and perfectly competitive quantities appear in the long run. The reason is that quantities closer to the perfectly competitive one tend to perform the best in each market, whereas markets closer to the Cournot quantity perform better than other markets. This creates countervailing tendencies in the dynamics which balance out in the long run. When firms use IBA, all firms produce a quantity strictly between the Cournot and Walrasian outcome in the unique stochastically stable state. This quantity decreases in the number of markets and increases in the number of firms per market. The outcome also corresponds to the unique symmetric Nash equilibrium of a game in which firms are concerned about profits in both absolute and relative terms.

In the second informational setting, firms observe markets where the aggregate quantity is sufficiently close. The idea is that imitation intuitively makes more sense if the sampled firms are in a similar situation as oneself. If firms are not very concerned about differences in the aggregate, the model behaves as in the benchmark case. If firms are very sensitive about differences in the aggregate, markets tend to evolve independently and we obtain the perfectly competitive outcome. For intermediate cases the outcome becomes more competitive as firms become more sensitive to differences in the aggregate and thereby less willing to imitate across markets.

In the third informational setting the markets are arranged around a circle and firms observe some of the neighboring markets. This reflects a setting in which firms' geographical locations prevent them from observing all the remaining markets. The results from the benchmark case are relatively robust to this setting. This means that if firms use IBA, the outcome becomes less competitive as firms observe a larger set of markets. A conclusion that holds across these different settings is that more information tends to lead to less competitive outcomes. This should be contrasted with the conclusion of Huck et. al. (1999, 2000), that more information *about the firms in the own market* leads to more competitive results.

This chapter is related to the seminal paper of Vega-Redondo (1997) who showed that a single market always converges to the perfectly competitive outcome. Here it is shown that when there are several markets, less competitive outcomes are obtained. The chapter is closely related to Apesteguia et. al. (2007), who were the first to discover that information about firms in other markets affects the outcome in models of imitation dynamics. However, the setting considered in this paper differs from that of Apesteguia et. al. (2007) in several ways. Apesteguia et. al. (2007) consider linear demand, zero costs, a capacity constraint at the Walrasian quantity and markets that are remixed in each period. Here we only assume that demand is decreasing and concave, that the costs are increasing and convex and there is no capacity constraint and markets are not remixed. The conclusions obtained are also clearly different from those of Apesteguia et. al. (2007). The benchmark case of the present paper can be seen as thoroughly studying some of the dynamics first uncovered by Apesteguia et. al. (2007). This chapter is also related to the single market models with memory analyzed in Alós-Ferrer (2004) and Bergin and

Berghardt (2009). This relationship comes from the fact that the memory model can be seen as a multimarket model, in which the additional markets exist in the memories of the firms.

In the third chapter "Altruism and Local Interaction" I study the altruistic provision of a local public good when imitation is more local than the public good. *Altruism* refers to acts beyond the self-interest. There are several approaches to reconcile such acts with the rationality typical of many economic models. One is to modify the utility function (Becker, 1974 and 1981), or to consider infinitely repeated interaction (Fudenberg and Maskin 1986). Recently, an alternative explanation was proposed by Eshel, Samuelson and Shaked (1998). They show that if individuals decide whether to altruistically provide a local public good by imitating the closest neighbors (with whom the good is shared), then altruistic behavior can persist. The assumption that imitation is local is reasonable, since it requires knowledge of others' actions and payoffs, which we have only about a small subset of society's members. However, many externalities are relatively *less* local. For example, littering close to where we live affects people that we do not even know. In particular, it is likely to affect a larger set of individuals than the ones we learn from.

In this chapter, I therefore analyze a model where the public good is assumed to be less local than imitation. Individuals live on a circle and decide whether to provide a public good (to be an altruist), or not do it (and be an egoist). The public good is shared by an arbitrary number of neighbors. To make a decision, the individual observes the actions and payoffs of the two closest neighbors and imitates the action that generated the highest average payoff. The individuals also sometimes experiment, in which case they choose an action randomly. In the absence of experimentation, altruism can persist and coexist with egoistic behavior as long as the public good is non-global. With experimentation, altruism can persist in the long run, but only if the population is large. Hence, the conclusion is that the result of Eshel et. al. (1998), is relatively robust to settings where imitation is more local than the public good.

The intuition behind the result is that since the public good is local, altruists can group together and exclude egoists. In this way, altruists have mainly altruist neighbors, and egoists have mainly egoist neighbors. Therefore, the altruists have higher payoffs and tend to be imitated. As the public good becomes less local, altruists need to gather in larger groups in order to exclude the egoists. As the public good becomes nearly global, at most one altruist group can fit in the population. It turns out that such a constellation is very sensitive to experimentation and it is actually enough that a single altruist switches to egoism for the population to descend into egoism. This is why a large population is needed for altruism to persist when the public good becomes less local.

The paper is related to Jun and Sethi (2007), Matros (2008) and Mengel (2009). In Jun and Sethi (2007) and Matros (2008) the public good is shared with the same (possibly large) set of individuals that each individual imitates. In these papers the local interaction structure is not as stylized as in Eshel et. al. (1998). But the assumption that the public good is shared with the individuals that are imitated is unlikely to hold when the good is almost global. Mengel (2009) instead assumes that imitation is *less* local than the public good. She finds that altruism will not survive in the long run. However, in many cases it is likely that precisely the opposite is true. The external effects of our actions often affect a larger number of individuals than the ones we learn from.

The fourth chapter studies strategic transmission of verifiable information when there are communication costs. Strategic communication is an important aspect of many economic situations in which some party is privately informed. In an influential paper Milgrom (1981) showed that if information can be *verifiably* communicated without costs and preferences satisfy a monotonicity requirement, then all private information is revealed. This is known as the *unraveling result.* However, reporting information is many times costly. It often requires careful and detailed explaining based on facts. It may also involve costly certification by accredited institutions. At the same time, the existing literature indicates that unraveling is sensitive to the assumption of costless reporting. (Jovanovic 1982, Verecchia 1983 and Cheong and Kim 2004).

In this work I study a model, in which, in contrast to previous literature, the cost of verifiably reporting private information *continuously* increases in the precision of the report. The conclusion is, contrary to previous work, that unraveling is rather robust to costly reporting. The model consists of a privately informed *Sender* and an uninformed *Receiver.* The Sender chooses which information to include in a report, which is given to the Receiver, who responds by choosing an action. The cost of the report increases continuously in the amount of information it contains. The Sender wants the Receiver to believe that he is of as high type as possible. It is shown that a separating equilibrium always exists, regardless of the reporting costs. The intuition is that the costs function as a signaling device. When the costs are high, a high Sender type can discourage lower types from mimicking the report through the costs. A combination of disclosure of information and costly signaling can always accomplish full separation. It is further shown that all separating equilibria are payoff equivalent. When the costs are low, all equilibria are separating. The reason is that with low reporting costs, high sender types can always break out of any pooling equilibrium by disclosing their true type. When the costs are high, a pooling equilibrium emerges. The intuition is that with high costs, Receiver scepticism makes it too expensive for the Sender to disclose his true type.

The reason that Jovanovic (1982), Verecchia (1983) and Cheong and Kim (2004) reach a different conclusion, is that the Sender is restricted to either report all his information, or none of it. This means that there are always some Sender types that will not disclose their information and if the cost is too high, the sender never reports any information. Here, the Sender can instead continuously decide how much information to report. When the reporting costs are arbitrarily high, the sender can thus report a small amount of information without incurring too high costs.

Finally, an extension is introduced in which the receiver must make a costly effort to access the information in the report. The idea is that it requires both time and effort to understand a report. I assume that the receiver can condition her reading decision on the appearance of the report, that the appearance is related to the amount of information in the report and that appearances can be manipulated at a cost. Two classes of separating equilibria are characterized. In *non-reading equilibria*, the receiver never reads any report. Depending on the reporting and manipulation costs, the sender either does not manipulate the appearance, or makes it look more precise than what it is. In *reading equilibria*, the receiver reads all the reports. Since both equilibria are separating, the receiver prefers the former, in which she incurs no costs.

This chapter is related to Mathis (2008), who considers communication in terms of *partially* verifiable information. Mathis' (2008) approach to partial verifiability can be treated as a model of costly reporting in which reports are either costless or arbitrarily costly. I.e., the reporting costs are discontinuous. In the present paper it may also be arbitrarily costly, to report all private information. However, in contrast to Mathis (2008), a separating equilibrium exists. The difference arises due to the continuity of the reporting costs in the present work. Hence, when treating partial verifiability in terms of costly reporting, the continuity of the costs has an important impact on the outcome. Another related paper is Kartik's (2009) work on costly lying. In Kartik's (2009) model, the sender can provide false information at a cost. However, in contrast to the present paper, the cost of the report is unrelated to its precision. Kartik (2009) finds that full separation is impossible and instead characterizes equilibria in which low types separate and high types pool.

Bibliography

[1] Akerlof, G. (1970). The Market for Lemons. *Quarterly Journal of Economics 84*, 488-500.

[2] Alós-Ferrer, C. (2004). Cournot vs Walras in Dynamic Oligopolies with Memory*. International Journal of Industrial Organization 22,* 193-217.

[3] Alos-Ferrer, C. and Schlag, K. (2009). Imitation and learning. Ch. 11, The Handbook of Rational and Social Choice, Anand, P.; Pattanaik, P.; Puppe, C., (Eds.) pp. 271-298(28), Oxford University Press.

[4] Apesteguia, J., Huck, S. and Oechssler, J. (2007). Imitation: Theory and Experimental Evidence. *Journal of Economic Theory, 135 ,* 217-235.

[5] Apesteguia, J., Huck, S., Oechssler, J. and Weidenholzer, S. (2010). Imitation and the Evolution of Walrasian Behavior: Theoretically Fragile but Behaviorally Robust. *Journal of Economic Theory 145*, 1603-1617.

[6] Bandura, A. (1977). Social Learning Theory. New York: General Learning Press.

[7] Becker, G.S. (1974). A theory of social interactions. *Journal of Political Economy 82*, 1063-1093.

[8] Becker, G.S. (1981). A Treatise on the Family. Cambridge: Harvard University Press.

[9] Bergin, J. and Bernhardt, D. (2009). Cooperation through Imitation. *Games and Economic Behavior 67*, 376-388.

[10] Cheong I. and Kim J.-Y. (2004). Costly Information Disclosure in Oligopoly. *The Journal of Industrial Economics 51*, 121-132.

[11] Conlisk, J. (1996). Why Bounded Rationality?*. Journal of Economic Literature 34*, 669-700.

[12] Ellison, G. and Fudenberg, D. (1993). Rules of thumb for social learning. *Journal of Political Economy 101*, 612-643.

[13] Ellison, G. and Fudenberg, D. (1995). Word of mouth communication and social learning. *Quarterly Journal of Economics 110*, 93-125.

[14] Eshel I., Samuelson L. and Shaked A. (1998). Altruists, Egoists and Hooligans in a Local Interaction Model. *American Economic Review 88*, 157-179.

[15] Festinger, L.(1954). A Theory of Social Comparison Processes. *Human Relations 7*, 117-140.

[16] Freidlin, M. and Wentzell, A.D. (1984). "Random Perturbations of Dynamic Systems". New-York: Springer-Verlag.

[17] Fudenberg, D. and Levine, D. (1998). The theory of Learning in Games. Cambridge MA: MIT press.

[18] Fudenberg, D. and Maskin, E. (1986). The Folk Theorem in Infinitely Repeated Games with Discounting or with Incomplete Information. *Econometrica 54*, 533-554.

[19] Gigerenzer, G., Todd, P.M. and the ABC Research Group. (1999). Simple Heuristics that Make us Smart. New York. New York: Oxford University Press.

[20] Huck, S., Normann, H.T, and Oechssler, J. (1999). Learning in Cournot Oligopoly: an Experiment. *Economic Journal 109,* 80-95.

[21] Huck, S., Normann, H.T, and Oechssler, J. (2000). Does Information about Competitors Actions Increase or Decrease Competition in Experimental Oligopoly Markets?. *International Journal of Industrial Organization 18*, 39-57.

[22] Jovanovic, B. (1982). Truthful Disclosure of Information. *The Bell Journal of Economics* 13, 36-44.

[23] Jun, T. and Sethi, R. (2007). *Journal of Evolutionary Economics 17,* 623-646.

[24] Kandori, M., Mailath, G. and Rob, R. (1993). Learning, Mutation and Long Run Equilibria in Games. *Econometrica 61*, 29-56.

[25] Kartik, N. (2009). Strategic Communication with Lying Costs. *The Review of Economic Studies 76*, 1359-1395.

[26] Mathis, J. (2008). Full Revelation of Information in Sender-Receiver Games of Persuasion. *Journal of Economic Theory 143*, 571-584.

[27] Matros, A. 2008. Altruistic Versus Rational Behavior in a Public Good Game. University of Pittsburgh, Working Paper.

[28] Mengel, F. (2009). Conformism and Cooperation in a Local Interaction Model. *Journal of Evolutionary Economics 19*, 397-415.

[29] Milgrom, P. (1981). Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics, 12*, 380-391.

[30] Offerman, T. and Schotter, A. (2009). Imitation and Luck: An Experimental Study of Social Sampling. *Games and Economic Behavior 65*, 461-502.

[31] Oyarzun, C. and Ruf, J. (2009). Monotone imitation. *Economic Theory 41*, 411-441.

[32] Pingle, M. and Day, R. (1996). Modes of Economizing Behavior: Experimental Evidence*. Journal of Economic Behavior and Organization 29,* 191-209.

[33] Sandholm, W. (2010). Population Games and Evolutionary Dynamics. Cambridge MA: MIT press.

[34] Schlag, K. (1998) Why imitate, and if so, how? A bounded rational approach to multi-armed bandits. *Journal of Economic Theory 78*, 130-156.

[35] Schlag, K. (1999). Which one should I imitate. *Journal of Mathematical Economics 31*, 493-522.

[36] Suls J., Martin R., and Wheeler L. 2000. Three Kinds of Opinion Comparison: the Triadic Model. *Personality and Social Psychology Review 4*, 219-237.

[37] Vega-Redondo, F. (1997). *The Evolution of Walrasian Behavior. Econometrica 65*, 375-384.

[38] Verecchia, R. (1983). Discretionary Disclosure. *Journal of Accounting and Economics 5*, 179-194.

[39] Weibull, J. (1995). Evolutionary Game Theory. Cambridge MA: MIT press.

[40] Young, P. (1993). The Evolution of Conventions. *Econometrica 61*, 57-84.

[41] Young, P. (1998). Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton: Princeton University Press.

Chapter 1

Social Comparison-based Learning in a Heterogeneous Population

1.1 Introduction

In this paper we study decision making processes when little information is available about the outcome an individual may obtain when choosing among different actions. While individuals may use information from their own previous experience, they may also use information obtained by other individuals. However, in many situations it is unlikely that the choices of one individual would lead to similar outcomes for another. For instance, a person whose height is less than 1.75 meters, can hardly be expected to play basketball as well as Kobe Bryant or Pau Gasol of the Los Angeles Lakers, whose heights are 1.98 and 2.13, respectively. Similarly, a monopolist with low production costs maximizes profits charging a low price, while another one with higher costs may be better off with a higher price. Even if individuals face cognitive and informational restrictions, they may be aware of these differences and take them into account in their decision processes. This heterogeneity and how individuals deal with it has received little attention in the literature on social learning and almost no attention within the literature of boundedly rational social learning. This paper is a first attempt to fill this gap.

We approach the problem by incorporating a notion of *social comparison* (e.g., Festinger 1954, Suls, Martin and Wheeler 2000) into a simple model of boundedly rational social learning. The theory of social comparison argues that in the absence of objective information, individuals have a drive to assess their skills, abilities and attributes *relative* to other individuals[1]. Furthermore, individuals use these assessments to evaluate their potential performance relative to others at different tasks. We show that when comparisons like these are part of individuals' decision processes in the presence of heterogeneity, outcomes are obtained that are qualitatively different from those identified by the previous literature. In particular, there is a strong tendency to obtain outcomes that favor some, but hurt others. Indeed, it is not unlikely to reach a state in which a fraction of all individuals always makes a good choice, whereas all others consistently make a bad one. More specifically, final outcomes tend to favor (1) the type of individuals that constitutes a majority of the population and (2) individuals that benefit the most from their optimal choice over their suboptimal choice.

We consider a population of individuals who repeatedly make a choice between two actions, which are the same across all individuals and time periods. As a consequence of her choice, each individual receives a payoff according to some unknown payoff distribution. These distributions may vary across individuals and may be interpreted as being determined by the individual's type. Yet, we assume that individuals do not know their own type, nor the type of any other individual in the population. The structure of information we consider is similar to the canonical setup in which decision makers are able to observe the experience of others (e.g., Fudenberg and Ellison, 1993, 1995; Vega-Redondo 1997; Schlag 1998, 1999). In particular, in each time period each individual observes the most recent action and payoff of an individual who is selected at random from the population according to an exogenous probability distribution. She also observes her own most recent action and consequent payoff. In contrast to the existing literature, we assume that each

[1] In economics, some papers have also studied social comparisons (e.g., Santos-Pinto and Sobel 2005), however this literature has focused on explaining observed over confidence and positive self assessment, whereas here we assume that decision makers compare themselves to others to assess what payoffs to expect with observed actions.

individual also perceives a random *comparison-signal*. This signal is informative about the difference in expected payoff of the individual and the observed individual with respect to the sampled action. We assume the comparison-signal to be unbiased, i.e., individuals do not make systematic mistakes in their assessment of the difference in expected payoff. The comparison signal can be understood as the individual's assessment of how much better or worse she would do than the observed individual, if she were to choose the same action as him.

Individuals use all this information to update their choice. The observed actions, obtained payoffs, and the perceived comparison-signal are mapped to the probability of choosing each action through a function called the *decision rule*. We impose a standard "must see" condition (e.g. Ellison and Fudenberg 1995, Cubitt and Sugden 1998), such that individuals consider switching action only when they observe an action different from the one that they are currently using. I.e., if an individual samples someone who is currently using the same action as herself, she does not switch.

Our analysis is confined to decision rules with which an individual who observes two different actions is always more likely to choose the one with the highest expected payoff for her. We call this class of decision rules *payoff-ordering*. In a preliminary result (Proposition 1.1) we show that for payoff-ordering decision rules, the updated probability of choosing each action depends on the difference between two terms: (1) the sum of the observed individual's payoff and the comparison signal and (2) the own obtained payoff. Intuitively, with payoff-ordering decision rules the probability of playing each action depends on the perceived difference in expected payoffs of the two actions. Our assumption that the comparison signal is unbiased implies that the perceived expected payoff of the sampled action is an unbiased estimator of the objective expected payoff of that action. Hence, this leads to a higher expected probability of choosing an action when its expected payoff is higher. More importantly, the comparison-signal assures that heterogeneity *per se* does not mislead individuals in their assessment of the expected payoff that they themselves would receive upon switching to the observed individual's action.

The focus of this paper is on the dynamics of the choices of individuals in a heterogenous population in which all of them update the probabilities of playing each action using a payoff-ordering decision rule. We consider a population with two types of individuals, A and B, who may choose between two actions, a and b. The payoff distributions associated with each action are the same for all individuals of the same type and the expected payoff of action a (b) is greater than the expected payoff of action b (a) for all type $A(B)$ individuals. We first analyze the case in which each individual observes each of the remaining individuals with uniform probabilities, which we refer to as *uniform sampling*. Our main results (Propositions 1.2 and 1.3) show that the choices of the population converge to a rest point determined by two factors: (1) the fraction of types in the population and (2) each type's difference in expected payoffs across actions, i.e. the *risk-sensitivity* of each type's decision problem (e.g., Persico 2000). If the two types' risk-sensitivity is about the same, then the relative fractions of types in the population determine the long run prediction of the dynamics. In particular, if the fractions of the two types are different enough, then the whole population converges to the action that provides the highest payoff for the majority type and hence, to the action with the lowest payoff for the minority. Intuitively, in order for an individual to switch to her optimal action she must observe someone who is using it. At the same time, each individual is more likely to observe someone from the majority type, who, due to the payoff ordering decision rule, is likely to be choosing his optimal action. This leads the optimal action of the larger type to propagate faster than the one of the smaller type. Eventually the entire population converges to the majority type's optimal action. Hence, the minority may end up choosing the wrong action in the long run, despite using a rule with which they are more likely to make the right choice among the observed actions.

The action to which the population converges is also determined by the risk-sensitivity of each type's decision problem. If the fractions of types are about the same, and the difference in risk-sensitivity across types is large enough, then the whole population converges to the optimal action of the most risk-sensitive type. In other words, the population converges to the optimal action of the type that benefits the most from moving to its optimal action. Intuitively, since choices are driven by payoff-ordering decision rules, and the relative fractions of types are about the same, the optimal action of the type that benefits the most from choosing optimally propagates faster than the other action. This eventually leads the system to converge to the action that provides the highest payoff for the most risk-sensitive type.

If the fractions of types in the population are similar and the risk-sensitivities of both types decision problems are about the same, then the population converges to a rest point in which a fraction strictly between zero and one of each type chooses its optimal action. Moreover, most individuals of at least one of the types choose their optimal action. While not all individuals make the right choice, payoff ordering assures that at least one type of individuals perform better than if they were to choose randomly. In this case, the balanced composition of the population make both actions propagate in a more balanced way, which eventually leads the population to a more balanced final state. However, in such a final state individuals from time to time sample an individual choosing a different action and this implies that some individuals of each type choose their suboptimal action.

We show (Proposition 1.4) that in all the asymptotically stable rest points of the system the average expected payoff of the population is greater than the average expected payoff if all individuals randomize their choices. In other words, if one type is better off in a stable rest point and the other type is worse off, then the gains of the type that is better off more than offset the losses of the other type. The intuition behind this result is simple. The system tends to converge to a state that favors the type that is either larger or face the most risk-sensitive decision problem, and this has a positive impact on the average expected payoff of the population.

Individuals may have a stronger (or weaker) tendency to observe individuals that are similar to them. We take this into account by allowing *homophilous* as well as *heterophilous* sampling, i.e. a bias toward observing the own and the other type, respectively. We refer to this as *biased sampling*. Homophilic tendencies are widely documented (see Currarini, Jackson and Pin 2009 and the references therein) and may be due to segregation or preferences for having friends that are similar to you. We show (Proposition 1.5) that the predictions are qualitatively similar to the case of uniform sampling, with some important qualifications. Each type is favored by being more homophilous. The reason is that this creates a stronger tendency for an individual to observe someone with the same optimal action, which in turn makes her more likely to observe an individual that is currently using her optimal action. On the other hand, each type is negatively affected by an increase in the other type's homophily, since this makes it more likely that the individuals of the opposite type are using their optimal action. Finally, as both types become more homophilous, the positive effects of the bias in sampling dominates and the population outcome improves. Thus, in our model segregation leads to better outcomes.

Finally, we discuss how our model can be applied to analyze the *diffusion of innovations*. There is an important literature that takes note of the fact that many innovations are not adopted instantaneously and study the process through which innovation diffuse in society (see, e.g, Geroski 2000 and Young 2009). Our model can be used to analyze the diffusion of innovation through social learning, when the innovations is beneficial for a fraction of the population, whereas the remainder is better off with the status quo. We discuss some related models and show that our model tends to generate the S-curves typical of many empirical studies of the diffusion of innovations[2]. We argue that our model is consistent with some features of the diffusion of hybrid corn in Kenya found by Suri (2011). In particular, Suri (2011) finds less than full final adoption, heterogeneities in returns to adoption and equilibrium switching behavior. All of these features are consistent with our model.

This paper is related to the seminal papers of Ellison and Fudenberg (1993,

[2]I.e., the rate of adoption is first slow, later fast and finally levels out.

1995), Vega-Redondo (1997), and Schlag (1998), which model individuals as having feedback from their own and others' choices, and using cognitively simple rules to choose their actions (for a thorough survey, see Alos-Ferrer and Schlag 2009). In the analysis of decision problems with a homogeneous population, relatively simple imitation rules allow most individuals to choose the optimal action in the long run. An important issue, not addressed by this literature, is the possibility that individuals are heterogeneous, and that being able to infer something about this heterogeneity may affect the choices and population outcome. A remarkable exception is Ellison and Fudenberg (1993), in which individuals only observe neighbors who, in most of the cases, have the same optimal actions as them. In fact, the smaller the "window" of neighbors an individual interacts with, the closer is the system to convergence to the optimal choice for the whole population in the long run. In contrast, here individuals are able to extract relevant information also from the experiences of individuals that are different from them, and observe all individuals with positive probability. This creates qualitatively different dynamics and e.g. allows for the possibility that minorities end up choosing a suboptimal action.[3]

In Section 2 we provide the framework and a characterization of payoff-ordering decision rules. Section 3 describes the dynamics of the two-type population. In Section 4 apply our model to the analysis of technology adoption. Section 5 explores the implications of biased sampling. Section 6 concludes.

1.2 Framework

We analyze a continuum population of individuals, denoted by W. Our model is intended to capture a situation in which the expected outcome of different choices

[3]Our paper also relates to the growing literature that studies empirically the role of social interactions in shaping individual choices, e.g. Manski 2000,Munshi 2004, A. Sorensen 2006, Bayer, Ross and Topa 2008. This literature has identified two factors that are fundamental to understand the role of social interactions on individual decisions, referred to as "endogenous-interaction effects" and "correlated-effects." Our model provides a description of how these effects may arise as a consequence of social learning. Endogenous interaction effects arise in our model since individuals are influenced by the population state through the sampling procedure. Correlated effects arise since individuals of the same type make similar choices.

may be different across individuals. For simplicity, we consider a population composed by two types of individuals, denoted by A and B, i.e., $W = A \cup B$. Let $\tau(i) \in T := \{A, B\}$ denote the type of individual $i \in W$. Each individual chooses an action $c \in S = \{a, b\}$. The action chosen by an individual i, denoted by $c(i)$, yields a payoff $x \in [0, 1]$, i.e., we assume there is a lower and upper bound for payoffs. The distribution of the payoff that an individual $i \in \tau$ receives when she plays action $c(i) = c$ is the same for all the individuals of that type, and it is denoted by $F_{\tau c}$, for all $\tau \in T$ and $c \in S$. The corresponding expected value is denoted by $\pi_{\tau c} := \int x dF_{\tau c}(x)$. The probability measure of every measurable set $X \subseteq [0; 1]$ under the distribution $F_{\tau c}$ is denoted by $\mu_{\tau c}(X)$ and if X is a singleton $\{x\}$ then, its probability is denoted simply by $\mu_{\tau c}(x)$. Notice that we use c to denote the mapping from $i \in W$ to her chosen action at any time, and also to denote a typical action in S (we may also use d and e to denote generic actions). In the same manner, a generic type is often denoted by τ or τ'.

From time to time, each individual i observes another individual $j \neq i$ in the population and finds out the action he played and the payoff he obtained. The sampling procedure is governed by exogenously given probabilities, denoted by $(p_{\tau,\tau'})_{\tau,\tau' \in T}$. In other words the probability that i observes j is determined only by $\tau(i)$ and $\tau(j)$. At the same time the individual realizes that she may be of a different type than the observed individual and, hence the payoff distribution she faces when playing an action may be different from that of the observed individual. In particular we assume that whenever a type τ individual observes a type τ' individual who has chosen c, she perceives a random variable $\delta_{\tau\tau' c}$ which takes values in $[-1, 1]$ and that we refer to as the *comparison-signal.* This random variable is informative about her relative performance if she were to choose c, compared to the performance of the observed individual when he chooses c. This comparison-signal is assumed to be unbiased in that $\int \delta dF_{\delta_{\tau\tau' c}}(\delta) = \pi_{\tau c} - \pi_{\tau' c}$. In other words, the expected value of $\delta_{\tau\tau' c}$ is positive (negative) when the individual who perceives it would do better (worse) with c than the observed individual. For instance, if i perceives $\delta > 0$ when she samples j who chose c, this may be interpreted as a judgement "I would do better than what j did if I choose c." The payoff of all individuals and perceived comparison-signals are all assumed to be mutually independent. The profile of distributions $F := (F_{\tau c}, F_{\delta_{\tau\tau' c}})_{\tau,\tau' \in T, c \in S}$ is called the *environment* of the decision problem and it is assumed to be unknown to the decision makers.

Individuals choose actions according to a *decision rule*, which is a mapping from the observed actions and corresponding payoffs, and the perceived comparison-signal, to the probability of choosing each action the next time. We denote this decision rule by L, thus $L : S \times [0,1] \times S \times [0,1] \times [-1,1] \to \Delta(S)$. Here, $L(c,x,d,y,\delta)(e)$ is the updated probability that an individual chooses $e \in S$, given that she chose c, obtained the payoff x, observed an individual who chose d and obtained the payoff y, and perceived the comparison-signal δ. As is common in the literature (e.g. Ellison and Fudenberg 1995, Cubitt and Sugden 1998), we assume $L(c,x,c,y,\delta)(c) = 1$ for all $c \in S$. This means that individuals only contemplate switching their action when they observe an action different from the one that they are currently using. This assumption is motivated in several manners (see Ellison and Fudenberg 1995 and the references therein).

For an individual $i \in \tau$, with decision rule L, given that she chooses action c and observes another individual $j \in \tau'$ who chooses action d, the expected probability that she chooses action e the next time she makes a choice is denoted by $L^e_{cd}(\tau,\tau')$, i.e.,

$$L^e_{cd}(\tau,\tau') := \int\int\int L(c,x,d,y,\delta)(e)dF_{\tau c}(x)dF_{\tau' d}(y)dF_{\delta_{\tau\tau' d}}(\delta).$$

In other words, $L^e_{cd}(\tau,\tau')$ captures the probability of choosing e, ex-ante to the draws from the payoff and comparison-signal distributions. This expression is important in the analysis of the decision rules that we focus on below.

It is important to observe that individuals are not assumed to know their types or, in other words, they do not know the payoff distribution of the different actions they may choose. They know neither the number of types in the population nor the distributions of comparison-signals. The only information they have is their own choice and payoff, and another individual's choice and payoff. Furthermore, individuals have a basic understanding of the heterogeneity in the population, incorporated in the comparison-signal they perceive. However, in spite of these severe information restrictions, there are simple rules of decision which allow individuals to often make good decisions. In particular our analysis focuses on decision rules which satisfy the property that each individual is most likely to choose the action that provides her with the highest expected payoff whenever she observes two different actions.

Definition 1.1. *A decision rule L is payoff-ordering if for any environment F, and actions c and d, if $\pi_{\tau d} > (<)\, \pi_{\tau c}$, then, $L^d_{cd}(\tau, \tau') > (<)\, L^c_{cd}(\tau, \tau')$ for all $\tau, \tau' \in T$.*

It turns out that the characterization of the class of payoff-ordering decision rules is simple and intiuitive. These rules choose the action of the sampled individual with a probability that is an affine transformation of the observed payoffs and the comparison-signal.

Proposition 1.1. *L is payoff-ordering if and only if $L(c, x, d, y, \delta)(d) = 1/2 + \beta(y + \delta - x)$, with $\beta \in (0, 1/4]$ for $c \in S, d \neq c, x, y \in [0, 1]$, and $\delta \in [-1, 1]$.*

We provide here the proof for the 'if' statement. The argument for the 'only if' part is lengthy so we provide it in Appendix A.

Proof. For any $\tau, \tau' \in T$, and $c, d \in S$,

$$\begin{aligned} L^d_{cd}(\tau, \tau') &= \int\int\int (1/2 + \beta(y + \delta - x)) dF_{\tau c}(x) dF_{\tau' d}(y) dF_{\delta_{\tau\tau' d}}(\delta) \\ &= 1/2 + \beta(\pi_{\tau' d} + \pi_{\tau d} - \pi_{\tau' d} - \pi_{\tau c}) \\ &= 1/2 + \beta(\pi_{\tau d} - \pi_{\tau c}). \end{aligned}$$

Therefore, if $\pi_{\tau d} > \pi_{\tau c}$, then $L^d_{cd}(\tau, \tau') > 1/2$. □

Since the comparison-signal is unbiased, $y + \delta_{\tau\tau' d}$ is an unbiased estimator of the expected payoff of individual $i \in \tau$ if she chooses the action chosen by the sampled individual $j \in \tau'$. Therefore, $y + \delta_{\tau\tau' d} - x$ is an unbiased estimator of the difference between the expected payoff of individual i when she chooses $c(j)$ and when she chooses $c(i)$. Hence, since $\beta > 0$, in expected terms, the probability of choosing the action that provides her the greatest expected payoff is greater than the probability of choosing the action that provides her the smallest expected payoff. Notice also that since $\beta \in \left(0, \frac{1}{4}\right]$, and $\pi_{\tau d} - \pi_{\tau c} \leq 1$, by the proof of Proposition 1.1, we have that $L^d_{cd}(\tau, \tau') \in \left(\frac{1}{2}, \frac{3}{4}\right]$ when $\pi_{\tau d} > \pi_{\tau c}$.

Previous work on imitation in homogeneous populations has focused on either some arbitrary, yet appealing rules, or decision rules that satisfy some desirable properties. For instance, Schlag (1998) studies *improving* rules, which satisfy that the average expected payoff of the population is expected to be non-decreasing in time. However in our model, heterogeneity rules out this possibility (see Claim 8 in Appendix B). Yet, every payoff-ordering decision rule is improving when the population is homogeneous (see Claim 9 in Appendix B).

It would be desirable that the decision rule guaranteed individuals not only to be more likely to choose their optimal action, but that they also did so with high probability. Unfortunately, this is not possible in our model. Indeed, it can be shown that for every decision rule, the expected value of the updated probability of choosing the action with the highest payoff of two observed actions, is arbitrarily close to one half for some environment. The (omitted) argument follows directly from the proof of Proposition 1.1.

Finally, we notice that perceiving comparison-signals allows individuals to make choices which in expectation are independent of the observed individuals type, and depend only on his choice. This result follows from the assumption that the comparison-signal is unbiased; the formal argument for its proof follows directly from the proof of Proposition 1.1 and it is omitted.

Remark 1. *For all* $i, j, k \in W$, $c, d \in S$, $L^d_{cd}(\tau(i), \tau(j)) = L^d_{cd}(\tau(i), \tau(k))$.

This means that in expectation, individuals make equally good choices regardless of whether they observe an individual of the same or of the other type. In the next section we analyze the dynamics of choices of a heterogeneous population in which all individuals make decisions according to a payoff-ordering decision rule. As we show below, even though each time an individual makes a choice she is more likely to choose her best action, frequent interaction with individuals choosing their own optimal action may lead her to end up choosing the worst action for her. Before proceeding to the population dynamics, we specify the dynamic structure of the model.

Dynamic Structure of the Model: The model outlined above can in principle proceed in discrete as well as in continuous time. In a discrete time version of the model, in each period each individual observes her action and consequent payoff in the previous period, as well as those of an individual selected at random from the population and a comparison signal. She uses this information to make a choice and in the next period the same pattern is repeated. Here, for analytical convenience we choose to focus on continuous time. A set of results very similar to those presented below for the case of continuous time holds in discrete time. The main difference between the continuous and discrete time versions of the model, is that in discrete time all individuals make their choices at exactly the same instant in time, whereas with continuous time, choices are made at different points in time.

We will assume that each individual is equipped with an artefact known as a *Poisson alarm clock* (see, e.g., Sandholm 2010). Each time this clock rings an individual is given an opportunity to make a choice. All clocks are independent and the times between the rings of an individual agent's clock follow a rate ρ exponential distribution. By a well known result in probability theory, this means that the number of rings in any time period $[0, t]$ follow a Poisson distribution with mean ρt. In other words, with larger ρ, individuals revise their choices more frequently. When the clock of $i \in W$ rings, she observes her most recent choice and payoff as well as those of another individual in the population and a comparison signal, as described above. She uses this information and a payoff-ordering decision rule to make a new choice. Once the choice is made a payoff is drawn from the corresponding distribution. This payoff can be thought of as a payoff-rate, which gives an instantaneous payoff to i at each point in time and remains the same until the next time the clock rings and the individual again makes a choice. It can also be thought of as a once and for all payoff that is given to the individuals at the exact time that she makes her choice. A possible concrete interpretation of this timing is that an individual buys a product, e.g. a computer, which breaks after some random amount of time and when this happens, the individual makes a new choice. More generally, the timing implies that individuals periodically revise their choices and the points in time at which they do this is given by an exogenous probability distribution.[4]

1.3 Population dynamics

In this section we analyze the dynamics of choices when individuals make their choices according to a payoff-ordering decision rule. We assume that a is the optimal action for type A and that b is the optimal action for type B. The probability that any type A individual is matched with another type A individual is denoted by $\alpha_A \in [0,1]$ and the probability that any type B individual is matched with a type B individual is denoted by $\alpha_B \in [0,1]$. We denote the sizes of the populations of types A and B by $|A|$ and $|B|$, respectively

From Remark 1, for any actions a, b and c in S, $L^c_{ab}(\tau, \tau')$ does not depend on τ'. Let $L^c_{ab}(\tau)$ denote the expected probability that a type τ individual, with $\tau \in \{A, B\}$, chooses c given that she played a and was matched with an individual playing b. Let $p(t)$ be the fraction of A types in the population choosing a at time t and let $q(t)$ be the fraction of B types in the population choosing b in t. Since we assume that the populations of type A and type B individuals consist of a continuum of individuals, the dynamics that arise from the process of sampling and social learning are approximately deterministic. Let $\dot{p}$ and $\dot{q}$ be the time derivatives of $p(t)$ and $q(t)$, respectively. Let the rate of the Poisson alarm clock be $\rho = 1$. The evolution of the fractions of individuals choosing a and b in the populations of A types and B types, respectively, is described by the system of quadratic differential

[4] Note also, that if the different individuals' alarm clocks are perfectly correlated, we obtain a discrete time dynamics.

equations

$$\begin{aligned}
\dot{p}(p,q) &= \alpha_A p(1-p)(L^a_{ba}(A) - L^b_{ab}(A)) &(1.1)\\
&\quad +(1-\alpha_A)((1-p)(1-q)L^a_{ba}(A) - pqL^b_{ab}(A))\\
\dot{q}(q,p) &= \alpha_B q(1-q)(L^b_{ab}(B) - L^a_{ba}(B)) &(1.2)\\
&\quad +(1-\alpha_B)((1-q)(1-p)L^b_{ab}(B) - qpL^a_{ba}(B)),
\end{aligned}$$

where the time dependence of p and q has been omitted. The first term on the right-hand side of (1.1) reflects the net flow to action a of type A individuals as a result of sampling within the population of A types. The term $(1-\alpha_A)((1-p)(1-q)L^a_{ba}(A)$ gives the flow to action a in the type A population as a consequence of sampling type B individuals choosing action a. Finally, $-(1-\alpha_A)pqL^b_{ab}(A))$ reflects the flow to action b in the population of type A individuals through sampling type B individuals choosing b. An analogous interpretation applies to (1.2).

Let $L := L^a_{ba}(A)$ and $R := L^b_{ab}(B)$ and notice that for payoff-ordering decision rules $L, R \in (1/2, 3/4]$. Since payoff-ordering decision rules satisfy $L^b_{ab}(\tau) = 1 - L^a_{ba}(\tau)$ for $\tau \in \{A, B\}$, the system of difference equations (1.1)-(1.2) can be written as

$$\begin{aligned}
\dot{p}(p,q) &= \alpha_A p(1-p)(2L-1) + (1-\alpha_A)((1-p)(1-q)L - pq(1-L)) &(1.3)\\
\dot{q}(p,q) &= \alpha_B q(1-q)(2R-1) + (1-\alpha_B)((1-q)(1-p)R - qp(1-R)). &(1.4)
\end{aligned}$$

From equation (1.3) we see that $\dot{p}(p,q)$ is increasing in L. Intuitively, a higher L means that type A individuals are more likely to choose a both when they play a and observe an individual playing b and vice versa. $\dot{p}(p,q)$ is decreasing in q. Intuitively, as q increases there are fewer type B individuals choosing a, therefore, when a type A individual samples a type B individual, the probability that this individual has played b is greater. This makes it more likely for type A individuals to choose b. The effect of p on $\dot{p}$ is ambiguous. In both a heterogeneous and a homogeneous population $\dot{p}$ is a concave polynomial in p. In a homogeneous population (or an isolated populations with $\alpha_A = 1$), when p is very small there are too few type A individuals from whom to imitate a. On the other hand, when p is very large, just a few type A individuals are left to switch from b to a. As long as $0 < p < 1$, $\dot{p} > 0$ because the flows from a to b are more than compensated by flows in the opposite direction, since $L > 1/2$. However in a heterogeneous (non-isolated) population, $p < 1$ is compatible with $\dot{p} < 0$. This occurs for high values of p, when an important fraction of A types may be mislead by encounters with type B playing b which may cause $\dot{p} < 0$. Yet, in the heterogeneous population, and for each $q \in (0,1)$, for small enough values of p, the response of $\dot{p}$ to increases in p are greater than for higher values of p. An analogous reasoning applies to $\dot{q}$ and (1.4).

In this section we assume that the sampling probabilities are uniform and hence, the probability that any individual observes a type A individual is $\alpha := |A|/(|A| + |B|) \in (0,1)$ and the probability that any individual observes a type B individual is $1-\alpha$. In this case $\alpha_A = \alpha$ and $\alpha_B = (1-\alpha)$. We abandon this assumption in Section 5 where we allow sampling to be biased.

1.3.1 The Rest points of the System

Let the set of rest points of the system (1.3)-(1.4) be denoted by RP, i.e.,

$$RP := \{(p,q) \in [0,1]^2 : \dot{p}(p,q) = \dot{q}(q,p) = 0\}.$$

First we characterize the set RP. It is natural that $(0,1)$ and $(1,0)$ are rest points, since in this case all individuals of both types choose the same action. Our next result shows that for some values of α, L, and R there is a third rest point located in the interior of $[0,1]^2$. We denote this point by $(\overline{p}^*, \overline{q}^*)$, with $\overline{p}^* := \frac{L(\alpha(L+R-1)-(1-L)(2R-1))}{\alpha(2L-1)(L+R-1)}$ and $\overline{q}^* := \frac{R((1-\alpha)(L+R-1)-(1-R)(2L-1))}{(1-\alpha)(2R-1)(L+R-1)}$. This third rest point exists if and only if $\alpha \in (\underline{\alpha}, \overline{\alpha})$, with $\underline{\alpha} := \frac{(1-L)(2R-1)}{L+R-1}$ and $\overline{\alpha} := \frac{L(2R-1)}{L+R-1}$.

Proposition 1.2. *If* $\alpha \in (\underline{\alpha}, \overline{\alpha})$, *then* $RP = \{(0,1), (1,0), (\overline{p}^*, \overline{q}^*)\}$, *otherwise* $RP = \{(0,1), (1,0)\}$.

Proof. Notice that if $p = 1$, then $\dot{p} = 0$ if and only if $q = 0$. Correspondingly, if

$q = 1$, then $\dot{p} = 0$ if and only if $p = 0$. This implies that $\{(0,1),(1,0)\}$ are rest points.

From (1.3) and (1.4), the points (p, q) that satisfy $\dot{p} = \dot{q} = 0$ are those that satisfy both

$$q = \frac{\alpha p(1-p)(2L-1) + (1-\alpha)(1-p)L}{(1-\alpha)((1-p)L + p(1-L))}, \quad (1.5)$$

$$p = \frac{(1-\alpha)q(1-q)(2R-1) + \alpha(1-q)R}{\alpha((1-q)R + q(1-R))}. \quad (1.6)$$

Define the function $\overline{q} : [0,1] \to \mathbb{R}$ such that, for every $p \in [0,1]$, $q(p)$ is given by the right hand side of (1.5). Therefore, for any $p \in [0,1]$, $q(p)$ gives the point q such that $\Delta p(p,q) = 0$. Analogously, the function $\overline{p} : [0,1] \to \mathbb{R}$ can be defined using the right hand side of (1.6) to determine $p(q)$. $\overline{p}(\cdot)$ and $\overline{q}(\cdot)$ are strictly concave in $[0,1]$ and they pass through $(0,1)$ and $(1,0)$. If (p^*, q^*) is a rest point of the system, then $p^* = p(q(p^*))$ and $q^* = q(p^*)$ (or equivalently, $q^* = q(p(q^*))$ and $p^* = p(q^*)$). Hence, for (p^*, q^*) to be a rest point, p^* must satisfy

$$p(q(p^*)) = \frac{(1-\alpha)q(p^*)(1-q(p^*))(2R-1) + \alpha(1-q(p^*))R}{\alpha((1-q(p^*))R + q(p^*)(1-R))} = p^*.$$

This yields $p^* = \overline{p}^*$. For this to be a rest point not in $\{(0,1),(1,0)\}$ we also need $p^* \in (0,1)$. The inequality $\overline{p}^* > 0$ simplifies to $\alpha > \underline{\alpha}$ and $\overline{p}^* < 1$ simplifies to $\alpha < \overline{\alpha}$. In other words $\overline{p}^* \in (0,1)$ if and only if $\alpha \in (\underline{\alpha}, \overline{\alpha})$. Further, it is obtained that $q(p^*) = \overline{q}^*$ which can be shown to lie in $(0,1)$ if and only if $\alpha \in (\underline{\alpha}, \overline{\alpha})$. Hence, there is a rest point given by (p^*, q^*) if and only if $\alpha \in (\underline{\alpha}, \overline{\alpha})$, and thus $RP = \{(0,1),(1,0),(p^*,q^*)\}$.

If $\alpha \notin (\underline{\alpha}, \overline{\alpha})$ there is no rest point in $(0,1)\times(0,1)$ and thus, $RP = \{(0,1),(1,0)\}$.

□

The rest point (p^*, q^*) only exists for some parameter values of α, L, and R. Abusing notation, let the functions $\underline{\alpha} : (1/2, 3/4]^2 \to \mathbb{R}$ and $\overline{\alpha} : (1/2, 3/4]^2 \to \mathbb{R}$ be respectively defined by $\underline{\alpha}(L,R) := \frac{(1-L)(2R-1)}{L+R-1}$ and $\overline{\alpha}(L,R) := \frac{L(2R-1)}{L+R-1}$ for all $(L,R) \in (1/2, 3/4]^2$. For an internal rest point to exist we need $\alpha \in (\underline{\alpha}(L,R), \overline{\alpha}(L,R))$. It is instructive to analyze how $\underline{\alpha}$ and $\overline{\alpha}$ respond to changes in L and R. Let $\underline{\alpha}_i$, and $\overline{\alpha}_i$

denote the first derivative of $\underline{\alpha}$ and $\overline{\alpha}$, respectively, with respect to its i^{th} argument for $i = 1, 2$. Then, differentiating, we obtain that $\underline{\alpha}_1(L, R) = \frac{-(2R-1)R}{(L+R-1)^2} < 0$. The interpretation of this is as follows. As is shown below, if $\alpha < \underline{\alpha}$, then the fraction of the population of type A is small, thus action a propagates slowly and the dynamics of the system drive the population to $(0, 1)$, i.e., virtually every individual in the population will choose b. When L is larger, this is prevented for a greater range of values of α and, hence, $\underline{\alpha}$ is smaller. In contrast, when R is large, a small range of values of α prevents that the population converges to choose b and, hence, $\underline{\alpha}$ is increasing in R. This provides an interpretation for $\underline{\alpha}_2(L, R) = \frac{(2L-1)(1-L)}{(L+R-1)^2} > 0$. On the other hand, as it is also shown below, if $\alpha > \overline{\alpha}$, then the fraction of the population of type A is so large that action a propagates very fast and the dynamics of the system drive the population to $(1, 0)$, i.e., virtually every individual in the population will choose b. Then an analogous interpretation can be provided for the signs of the partial derivatives $\overline{\alpha}_1(L, R) = \frac{-(2R-1)(1-R)}{(L+R-1)^2} < 0$ and $\overline{\alpha}_2(L, R) = \frac{(2L-1)L}{(L+R-1)^2} > 0$.

Let $(\hat{p}(L, R, \alpha), \hat{q}(R, L, 1-\alpha))$ denote the internal rest point, i.e.,

$$\begin{aligned} \hat{p}(L, R, \alpha) &= \frac{L(\alpha(L+R-1) - (1-L)(2R-1))}{\alpha(2L-1)(L+R-1)} \\ \hat{q}(R, L, 1-\alpha)) &= \frac{R((1-\alpha)(L+R-1) - (1-R)(2L-1))}{(1-\alpha)(2R-1)(L+R-1)}. \end{aligned}$$

$\hat{p}(L, R, \alpha)$ and $\hat{q}(R, L, 1-\alpha)$ can be expressed in terms of $\underline{\alpha}(L, R)$ and $\overline{\alpha}(L, R)$ as

$$\begin{aligned} \hat{p}(L, R, \alpha) &= \frac{L}{2L-1} \frac{\alpha - \underline{\alpha}(L, R)}{\alpha} \\ \hat{q}(R, L, 1-\alpha) &= \frac{R}{2R-1} \frac{(1-\alpha) - (1-\overline{\alpha}(L, R))}{(1-\alpha)}. \end{aligned}$$

This reveals that given L and R, the internal rest point depends on the distance of α from $\underline{\alpha}(L, R)$ and $\overline{\alpha}(L, R)$.[5] Then, for $\alpha \in (\underline{\alpha}(L, R), \overline{\alpha}(L, R))$ we obtain the following comparative statics result for the interior rest points. Let $\widehat{p_i}$, denote the first derivative of $\hat{p}(L, R, \alpha)$ with respect to its i^{th} argument, for $i = 1, 2, 3$. Then, differentiating, we obtain the following result.

[5] Notice that $\hat{q}(R, L, 1-\alpha) = \hat{p}(R, L, 1-\alpha)$.

Remark 2. *$\widehat{p}_1$, $\widehat{p}_2$, and $\widehat{p}_3$ satisfy*[6]

$$\hat{p}_1(L,R,\alpha) = \frac{(1-R)\underline{\alpha}(L,R)}{\alpha(2L-1)} + \frac{(\overline{\alpha}(L,R)-\alpha)(L+R-1)}{\alpha(2L-1)^2} > 0$$

$$\hat{p}_2(L,R,\alpha) = -\frac{L(1-L)}{\alpha(L+R-1)^2} < 0$$

$$\hat{p}_3(L,R,\alpha) = \frac{\underline{\alpha}(L,R)}{\alpha^2(2L-1)} > 0.$$

The intuition of these results is slightly subtle. A greater L leads to a higher probability that a type A individual chooses a both when she currently chooses a and samples an individual choosing b and vice versa. In the interior rest point (p^*, q^*) an increase in p makes $\dot{p}(p,q) < 0$. As $\dot{p}(p,q)$ is increasing in L, this allows that $\dot{p}(p,q) = 0$ for a higher value of p. An analogous reasoning reveals that an increase in R causes that in an equilibrium with greater value for R, we have a greater value for q as well. Since an increase in the fraction of type B individuals choosing b causes a decrease in $\dot{p}(p,q)$, the new equilibrium requires a decrease in the fraction of type A individuals choosing a. A similar, but slightly more intricate (omitted) argument reveals why an increase in α causes the fraction of type A individuals to increase and the fraction of type B individuals to decrease.

1.3.2 Stability of the Rest points

Next we analyze the conditions for the different rest points to be stable. The notion of asymptotic stability requires the system to remain close and to converge to the rest point whenever the system starts sufficiently close to it (e.g., Hofbauer and Sigmund 1998). Formally a rest point (p^*, q^*) is *asymptotically stable* if for any $\varepsilon > 0$ there exists some $\delta \in (0, \varepsilon)$ such that if $||(p(t), q(t)) - (p^*, q^*)|| < \delta$, then $||(p(t'), q(t')) - (p^*, q^*)|| < \varepsilon$ for all $t' > t$, and (ii) there exists some $\delta > 0$ such that if $||(p(t), q(t)) - (p^*, q^*)|| < \delta$, then $\lim_{t' \to \infty}(p(t'), q(t')) = (p^*, q^*)$.[7]

[6] Since $\hat{q}(R, L, 1-\alpha) = \hat{p}(R, L, 1-\alpha)$, (omitted) analogous results hold for $\hat{q}(R, L, 1-\alpha)$.
[7] $|| \cdot ||$ stands for the euclidean norm of the corresponding vector dimension.

In the following proposition we characterize the stability properties of the different rest points of system (1.3)-(1.4) for virtually all the possible parameter values.[8] Furthermore, we show that in each case, the asymptotically stable rest point is a *global attractor*, i.e., the system converges to this point, regardless of the initial conditions (as long as the path does not start in a different rest point). The proof is provided in the appendix.

Proposition 1.3. *Suppose sampling is uniform and let* (p^*, q^*) *be the asymptotically stable point of the system. Then*

$$(p^*, q^*) = \begin{cases} (0,1) & \textit{if} \quad \alpha < \underline{\alpha}(L,R) \\ (\hat{p}(L,R,\alpha), \hat{q}(R,L,1-\alpha)) & \textit{if} \quad \underline{\alpha}(L,R) < \alpha < \overline{\alpha}(L,R) \\ (1,0) & \textit{if} \quad \alpha > \overline{\alpha}(L,R). \end{cases}$$

Furthermore, in each of these cases $\lim_{t\to\infty}(p(t), q(t)) = (p^*, q^*)$ *for all paths such that* $(p(0), q(0)) \notin RP$.

Whenever a rest point is a global attractor, it satisfies the second condition of asymptotic stability. However, in general, it does not necessarily satisfy the first part, i.e. it is possible that the dynamics move away from the rest point before eventually converging to it. Nevertheless, Proposition 1.3 reveals that a rest point of (1.3)-(1.4) is asymptotically stable if and only if it is a global attractor of the system.

Proposition 1.3 implies that if there is an internal rest point, it is the only asymptotically stable point and additionally it is a global attractor. Otherwise, either $(0,1)$ or $(1,0)$ is asymptotically stable and a global attractor. In particular, if $\alpha < \underline{\alpha}(L,R)$ we have that $(0,1)$ is asymptotically stable, whereas if $\alpha > \overline{\alpha}(L,R)$ we have that $(1,0)$ is asymptotically stable. Intuitively, if $\alpha < \underline{\alpha}(L,R)$, the fraction of the population of type A individuals is relatively small, so b propagates much faster. This eventually leads the whole population to choose b. Analogously, if $\alpha > \overline{\alpha}(L,R)$, the fraction of the population of type A individuals is relatively large,

[8] We only exclude $\alpha = \underline{\alpha}$ and $\alpha = \bar{\alpha}$. The reason for this is that the techniques that we use in the proof of the following result applies only to hyperbolic rest points (see the proof of Proposition 1.3) and at $\alpha = \underline{\alpha}$ and $\alpha = \bar{\alpha}$ some of the rest points are not hyperbolic.

so a propagates much faster and eventually, the whole population ends up choosing a. Finally, if the fraction of the population of type A is neither large enough nor small enough, then neither a nor b fully dominate and $(\hat{p}(L,R,\alpha),\hat{q}(R,L,1-\alpha))$ is the asymptotically stable rest point. As we show below, in this rest point, at least one of the two types are better off than simply choosing their actions randomly.

1.3.3 Phase Diagrams

In this section we derive the phase diagrams corresponding to the system (1.3)-(1.4). This is done by using equations (1.5)-(1.6) and observing that

$$\dot{p}(p,q) \lesseqqgtr 0 \Longleftrightarrow q \gtreqqless \frac{\alpha p(1-p)(2L-1)+(1-\alpha)(1-p)L}{(1-\alpha)((1-p)L+p(1-L))} \tag{1.7}$$

$$\dot{q}(q,p) \lesseqqgtr 0 \Longleftrightarrow p \gtreqqless \frac{(1-\alpha)q(1-q)(2R-1)+\alpha(1-q)R}{\alpha((1-q)R+q(1-R))}. \tag{1.8}$$

In other words, the direction in which the system (1.3)-(1.4) moves in the (p,q) space given some initial point (p_0,q_0) will depend on whether this point is above or below the graphs of the functions $p(q)$ and $q(p)$ defined in the proof of Proposition 1.2. There are three possibilities: (i) "$p(q)$ is always above $q(p)$," i.e., $p^{-1}(\overline{p}) > q(\overline{p})$ for all $\overline{p} \in (0,1)$, (ii) "$p(q)$ is always below $q(p)$," i.e., $q(\overline{p}) > p^{-1}(\overline{p})$ for all $\overline{p} \in (0,1)$, and (iii) "$p(q)$ and $q(p)$ intersect," i.e., $q(\overline{p}) = p^{-1}(\overline{p})$ for some $\overline{p} \in (0,1)$. In Appendix 3, we provide the formal analysis that reveals that these are all the possible cases.

The three cases are described in Figures 1, 2 and 3, respectively, for a representative parameter configuration corresponding to each case. In all three plots the solid line represents $p(q)$ and the dashed line represents $q(p)$. The parameter values used for the plots are (i) $L=3/5$, $R=3/4$ and $\alpha=1/2$, (ii) $L=3/4$, $R=3/5$ and $\alpha=1/2$ and (iii) $L=R=3/4$ and $\alpha=1/2$.

The phase diagrams give an idea about the direction in which the system moves given different initial points. In case (i), in which $p(q)$ is above $q(p)$, the phase diagram in Figure 1 suggests that the system moves in the direction of $(0,1)$ from any initial point other than $(1,0)$. This case corresponds to $\alpha < \underline{\alpha}(L,R)$. In this case, although the population is perfectly balanced, the two-armed bandit faced by type B individuals is more risk sensitive than the bandit faced by type A individuals because $L = 3/5 < R = 3/4$.

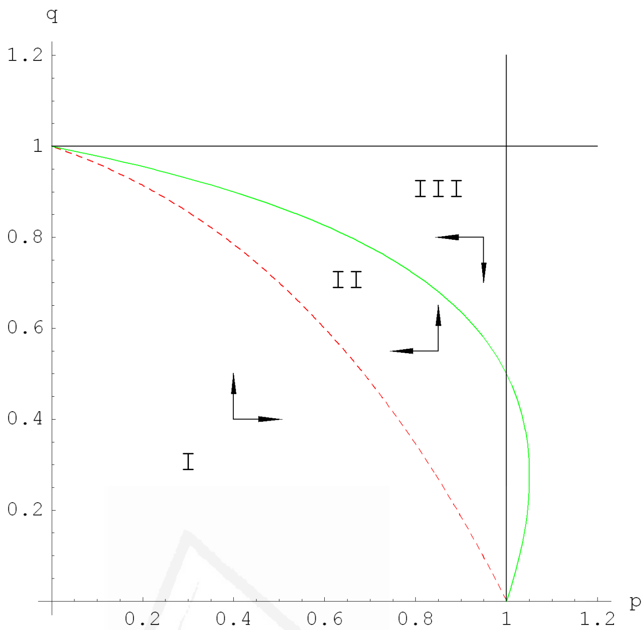

Figure 1: $p(q)$ is above $q(p)$: $L = 3/5$, $R = 3/4$ and $\alpha = 1/2$

Analogously, in case (ii), in which $q(p)$ is above $p(q)$, the phase diagram in Figure 2 suggests that the system moves towards $(1, 0)$ from any initial point different to $(0, 1)$. This case corresponds to $\alpha > \overline{\alpha}(L, R)$.

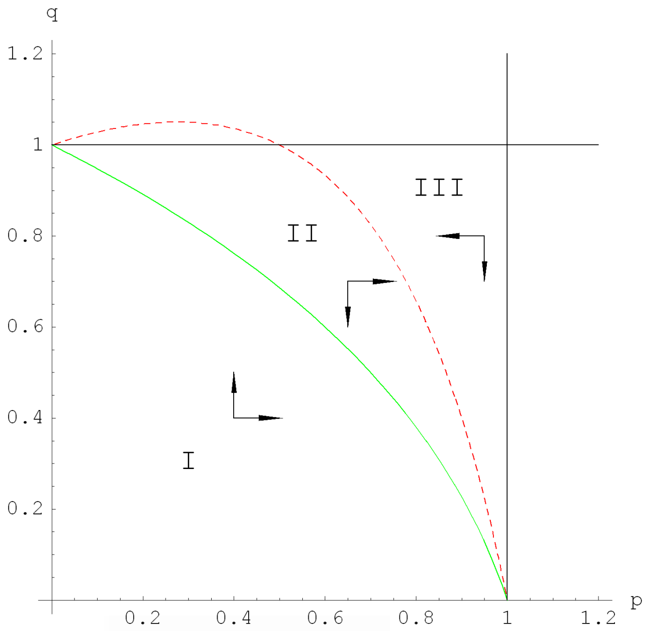

Figure 2: $q(p)$ is above $p(q)$, $L = 3/4$, $R = 3/5$ and $\alpha = 1/2$

Finally, in case (iii) (Figure 3) the system appears to be moving towards the point of intersection $(\hat{p}(L, R, \alpha), \hat{q}(R, L, 1-\alpha))$ from any initial point other than $(0, 1)$ and $(1, 0)$. It is easy to verify that in this case $\alpha \in (\underline{\alpha}(L, R), \overline{\alpha}(L, R))$.

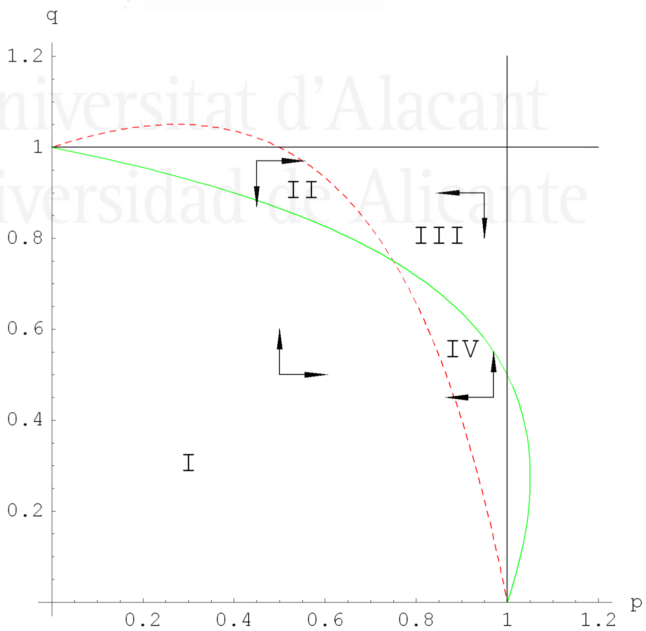

Figure 3: There is a point of intersection, $L = R = 3/4$ and $\alpha = 1/2$.

1.3.4 Average Expected payoffs

The minority's curse. The system converges to $(0,1),(1,0)$ or an interior point, depending on parameter values. Each population type benefits from being large relative to the other population type, in the sense that, if the fraction of type A (B) individuals is large enough, namely, $\alpha > \overline{\alpha}$ ($\alpha < \underline{\alpha}$), then the fraction of individuals of that type playing their optimal action, a (b), converges to one.

Since $\beta > 0$, both L and R are sensitive to the difference between the expected payoff between the two actions. This has the consequence that, for a relatively balanced population, if the difference between the expected payoff of actions a and b for the individuals in A (B) is large enough compared to the difference between the expected payoff of actions a and b for the type B (A) individuals, then the system converges to a state where all individuals choose a (b), the optimal action for type A (B) individuals.

It is noteworthy that the system predicts one of the extreme points for a relatively large set of parameter values. Thus, even though all individuals make their choices according to a payoff-ordering decision rule, the population dynamics may well lead one of the types to a state in which all of them choose the action that is optimal for the other type, but not optimal for them. The reason is that the forces leading the individuals of the other type towards their optimal action can be stronger and imitating them may overwhelm the effect of payoff ordering. In particular, imitation in a heterogenous population may be harmful for minorities. If α is sufficiently large, type B individuals would converge to b in the long run only if type A individuals are nearly indifferent between a and b. In other words, if a minority interacts with a majority and preferences over actions are the opposite, the majority will exert a negative influence on the minority leading them to choose the action that is not optimal. This contrast sharply with previous results obtained for homogenous population where these adverse effects do not arise.

The adverse effect for the minority is driven by the popularity of the suboptimal action, in the sense that more popular actions are chosen more frequently. A more popular action is more likely to be chosen since sampling an individual playing it is more likely. The larger the fraction of a type, the more influential their decisions will be for the population dynamics. For example, if a type B individual playing a is matched with another individual playing b, the probability that he will choose b is $R > 1/2$. However, the probability that he will choose b ex-ante to the outcome of the matching process will be $(\alpha(1-p)+(1-\alpha)q)\,R$ which is less than $1/2$ for many values of p and q. At the same time, the larger is α the greater is the weight of p in this probability. This contrasts with the manner in which popularity is introduced in Ellison and Fudenberg (1993). They assume that when individuals make their choices they are biased towards more popular actions. In our analysis, popularity weighting is not an exogenously imposed bias toward more popular actions, but it rather arises endogenously as a result of the sampling procedure.

The adverse effect over minorities however may not appear if the difference between the expected payoffs among actions for the majority is small. In fact, notice that $\lim_{R\to 1/2}\underline{\alpha}(L,R) = \lim_{R\to 1/2}\overline{\alpha}(L,R) = 0$ and $\lim_{L\to 1/2}\underline{\alpha}(L,R) = \lim_{L\to 1/2}\overline{\alpha}(L,R) = 1$. This implies that for small enough R (and for fixed α and L) the system converges to $(1,0)$, whereas for small enough L (and for fixed α and R) the system converges to $(0,1)$. In other words, even if, for instance, group A is much larger than group B, if A is close to indifferent between a and b the system will converge to a state in which the minority (B) chooses the right action, whereas A chooses the wrong one. On the other hand, since $\underline{\alpha}(L,R), \overline{\alpha}(L,R) \in (0,1)$, for any L, R there exist some α such that the system predicts either $(1,0)$ or $(0,1)$. This means that even if, for example, type A individuals are close to indifferent between a and b, if the population of individuals type A is sufficiently large relative to the population of individuals type B, the system will converge to a state in which all type A individuals choose the optimal action, while all type B individual choose their suboptimal action.

It is instructive to contrast the choices individuals make in an asymptotically stable rest point with simple random choice. When choice is random, on average, half of the individuals of each type choose the right action. We can thus consider imitation to be detrimental for a type of individuals whenever it leads to a population state in which less than half of its members chooses the right action in the asymptotically stable rest point. Imitation is detrimental for one of the types when the asymptotically stable rest point is either $(0,1)$ or $(1,0)$. At the same time, in an internal rest point imitation can only be detrimental for at most one of the population types, which follows since $p(q)$ and $q(p)$ are concave functions and both go through $(0,1)$ and $(1,0)$, which implies that if they intersect they do so above the line $q = 1-p$. In an internal rest point imitation is detrimental for type A individuals whenever $\hat{p}(L,R,\alpha) < 1/2$ which can be rewritten $\alpha < 2L\underline{\alpha}(L,R)$. Analogously, imitation is detrimental to type B individuals if $\alpha > 1-2R(1-\overline{\alpha}(L,R))$. Imitation is beneficial for both types of individuals when these two inequalities hold in the opposite direction.

The average expected payoff of the population. At most one of the two population types may be better off choosing randomly than when choices are lead by a payoff-ordering decision rule. However, if so, this type of individuals represents a small fraction of the population or has the smallest difference in expected payoffs between the optimal and suboptimal action. Let $\pi_X(a)$ ($\pi_X(b)$) be the expected payoff of an individual type $X \in \{A,B\}$ when she plays action a (b). The average expected payoff of the population in the state (p,q), denoted $W(p,q)$, is given by

$$W(p,q) = \alpha(p\pi_A(a) + (1-p)\pi_A(b)) + (1-\alpha)(q\pi_B(b) + (1-q)\pi_B(a)).$$

The average expected payoff of the population when all individuals choose randomly with uniform probability, denoted by W^{RC}, is given by

$$W^{RC} = \frac{1}{2}(\alpha(\pi_A(a) + \pi_A(b)) + (1-\alpha)(\pi_B(b) + \pi_B(a))).$$

We then obtain

Proposition 1.4. *$W(p^*,q^*) - W^{RC} > 0$ for any asymptotically stable rest point (p^*,q^*) of system (1.3)-(1.4).*

Proof. First, consider $\alpha > \overline{\alpha}(L,R)$. Then $(p^*,q^*) = (1,0)$ and $W(1,0) - W^{RC} = \frac{1}{2}(\alpha(\pi_A(a) - \pi_A(b)) - (1-\alpha)(\pi_B(b) - \pi_B(a))) = \frac{1}{4\beta}(\alpha(2L-1) - (1-\alpha)(2R -$

1)), which means that $W(1,0) - W^{RC} > 0$ if $\alpha > \frac{2R-1}{2(L+R)}$ and this holds since $\frac{2R-1}{2(L+R)} < \overline{\alpha}(L,R)$. An analogous argument holds if $\alpha < \underline{\alpha}(L,R)$. Suppose $\alpha \in (\underline{\alpha}(L,R), \overline{\alpha}(L,R))$. For simplicity assume $\beta = \frac{1}{4}$ (a similar argument holds if $\beta \in (0, \frac{1}{4})$. Then $W(\hat{p}(L,R,\alpha), \hat{q}(R,L,1-\alpha)) - W^{RC} = \alpha(\hat{p}(L,R,\alpha) - \frac{1}{2})(\pi_A(a) - \pi_A(b)) + (1-\alpha)(\hat{q}(R,L,1-\alpha) - \frac{1}{2})(\pi_B(b) - \pi_B(a)) = (2R-1)(2L-1) > 0$. □

This shows that in any interior rest point the imitation process leads to a higher average expected payoff, despite of the fact that it may be detrimental for one of the population types. In fact, the amount by which it increases welfare in an interior rest point is increasing in both L and R. In other, words, the gains over random choice of the individuals of the type that benefits from the imitation process always exceeds the loss incurred for the the type of individuals that is hurt by the imitation process. Note also that it is not necessarily the case that one of the types is worse as compared to random choice. For instance, if $\alpha = 1/2$, $L = 3/4$ and $R = 3/4$, then $(\hat{p}(L,R,\alpha), \hat{q}(R,L,1-\alpha)) = (3/4, 3/4)$, so the population converges to an asymptotically stable state in which 75% of individuals of each type choose their optimal action.

1.4 Application to the Diffusion of Innovations

There is a large literature studying how innovations are adopted and diffuse in society.[9] The basic observation is that innovations are rarely adopted instantaneously and it rather takes a long time for adoption to take place. A prominent stylized fact about slow diffusion is that the time path of the fraction of adopters follows an S-curve, i.e. adoption is first slow, then accelerates and finally slows down.[10] Much of the literature on diffusion of innovations aims to explain how such patterns of adoption may arise. In this section we illustrate how our model can used to analyze

[9] See for example, Rogers (1995), Geroski (2000), Young (2009), Conley and Udry (2010), and Suri (2011).

[10] For instance, see Ryan and Gross (1943), Griliches (1957), Dixon (1980), and Henrich (2001).

diffusion of innovations. We discuss some related models and derive some empirical predictions some of which are consistent with a number of empirical findings.

Suppose that a population is using a technology a, when a new technology b is introduced in the market. Suppose this technology provides, in expected value, higher payoffs than technology a for each type B individual. Yet, the expected payoffs of type A individuals is lower with the new technology than with the current technology a. Consider an initial state in which $q \simeq 0$ and $p = 1$, i.e. a small amount of type B individuals initially adopt the new technology b, and the remainder use the status quo technology a. All type A individuals use technology a in the initial state. We can then study how b is adopted by type A individuals, type B individuals, and the entire population. In particular, our model captures diffusion of an innovation which is optimal for a fraction of the population, whereas the remainder does better with the status quo. An example of such a situation is the diffusion of hybrid corn in Kenya examined by Suri (2011), where the returns to adoption depend on differences in infrastructure and other factors. Suri (2011) explains patterns of heterogeneity in adoption by considering rational heterogenous adopters who choose the technology that is optimal to them. Here we give a complementary explanation in which heterogeneity in adoption arises from a process of social learning and comparison when the payoffs of different technologies vary across individuals. The idea that social learning can influence diffusion of innovations is not new and in fact empirical evidence supports it. For example, Conley and Udry (2010) find evidence that farmers in Ghana are influenced by successful neighbors.[11] However, the impact of heterogeneity of payoffs across the population on the diffusion process, or how agents consider this heterogeneity in their decisions has received virtually no attention in the literature.

Contagion models. Diffusion arising from the type of social learning studied

[11] Foster and Rosenzweig (1995) reached a similar conclusion for Indian agriculture in an earlier study. Another example is Munshi (2003). He finds that individuals are less likely to learn from neighbors that differ from them in *unobserved characteristics*. This means that individuals do not have access to the information implicit in the similarity signal contemplated in our model. This information is based on observable characteristics and is what allows the extraction of relevant information even from the experience of an individual of a different type.

in this paper is closely related to a process of diffusion known as *contagion* (e.g., Geroski 2000, Young 2009). In models of contagion, individuals are assumed to adopt a technology when they observe or interact with others who have already adopted. Hence, adoption depends on the prevalence of the technology in the population. The model analyzed in this paper relates to a two-population model of contagion in which individuals may also *disadopt* at a positive rate, i.e. choose to abandon the new technology. To see this relationship we consider a relatively standard version of the model of contagion analyzed in Young (2009). In his analysis there is only one type of individuals in the population, say type B, and the dynamics of the fraction of the population who have adopted at time t, denoted by $q(t)$, are given by the equation of motion

$$\dot{q}(t) = (\lambda q(t) + \gamma)(1 - q(t)),\ \ \lambda, \gamma > 0, \tag{1.9}$$

where $q(t)$ is the fraction of current adopters and $\dot{q}(t)$ its time derivative.[12] The parameter λ is the (exogenous) rate at which individuals adopt the technology after hearing about it from another individual. The parameter γ measures adoption through external influences such as advertising. The term $q(t)(1 - q(t))$ can be thought of as the number of encounters between adopters and non-adopters.

The contagion model and equation (1.9) in particular bear some similarity to (1.3) and (1.4). This is not surprising, in our model individuals sample others from the population and mimic their action at some rate, which is similar to a process of contagion. To see more clearly the nature of this relationship, consider an extension of the model of contagion to the case of a two type population where $\gamma = 0$.[13] Let $1 - p$ and q be the fraction of agents who has adopted the new technology in populations A and B respectively. Suppose further that type A individuals who have not adopted the new technology b, adopt at rate λ_{Ab} whenever they have

[12] Young also considers a process of diffusion through social learning. However, in his model individuals are far more rational than in ours and update a prior belief as information accumulates. It turns out that our model of boundely rational social learning is more related to a process of contagion.

[13] A two-population model of contagion is discussed in Geroski (2000) who also provides a literature review on the contagion model.

encounters with other individuals of any type who have adopted b. Suppose also that disadoption may occur as well at rates λ_{Aa} when type A individuals who have adopted b have encounters with individuals who have not adopted. Similarly let the adoption and disadoption rates of type B individuals be denoted by λ_{Bb} and λ_{Ba}, respectively. Then, the resulting dynamics are given by

$$\begin{aligned} \dot{p} &= \alpha p(1-p)(\lambda_{Aa} - \lambda_{Ab}) + (1-\alpha)((1-q)(1-p)\lambda_{Aa} - pq\lambda_{Ab}) \quad (1.10) \\ \dot{q} &= (1-\alpha)q(1-q)(\lambda_{Bb} - \lambda_{Ba}) + \alpha((1-p)(1-q)\lambda_{Bb} - pq\lambda_{Ba}). \quad (1.11) \end{aligned}$$

The constants $\lambda_{Aa} - \lambda_{Ab}$, λ_{Aa}, and λ_{Ab} in (1.10) play the same role as $2L-1$, L, and $1-L$ in (1.3) (for (1.11) the corresponding analogies hold). Hence our social comparison-based learning model may be interpreted as a two-population contagion model with adoption and disadoption rates given by, respectively, $1-L$ and L, for type A individuals and R and $1-R$ for type B individuals.

Whereas disadoption would have little effect if incorporated in (1.9), the effect of disadoption in a heterogeneous two type population is important. For example, as shown in the previous sections, it is plausible that the population converges to a state in which only a fraction of it adopts the new technology. Oftentimes, disadoption seems plausible, in particular if the returns to the innovation are uncertain, and moreover it is worse than the status quo for a part of the population. Indeed, Young (2009) allows for disadoption in the analysis of the model of social learning and Suri (2011) finds evidence of disadoption and switching when returns to adoption may be negative. Additionally, by considering diffusion of innovations through the learning model considered here, we obtain dynamics governed by rates of adoption and disadoption that are related to economic fundamentals such as payoffs and the decision rules of individuals.

Adoption curves. The literature on technology diffusion has paid considerable attention to the *adoption curve*, i.e. the time path of the fraction of current users. The adoption curve of (1.9) has a standard S-shape that is often found in the empirical literature on technology adoption (e.g., Griliches 1950, Dixon 1980). Although finding explicit solutions to (1.3) and (1.4) is intractable, the adoption curves of our model are S-shaped for a range of typical parameter values. It is instructive to analyze the dynamics of the dicrete time version of (1.3) and (1.4).

We illustrate the dynamics of the model in two cases; one where full adoption occurs, and another where only partial adoption is observed. In Figure 4, the dash-dot line is the adoption curve for type B individuals, the dashed line corresponds to type A individuals and the solid line is for the entire population. On the left-hand side of Figure 4, $\alpha = 0.4$, whereas on the right-hand side $\alpha = 0.66$. In both parts $L = 0.6$ and $R = 0.7$ so, as before, a is better for type A individuals and b is better for type B individuals. In the left-hand side the entire population adopts the new technology, which means that 40% of the population ends up using the new technology in spite of being better off with the status quo. It can be seen that at any point in time type B individuals have a larger fraction of adopters than type A individuals. Eventually, most type B individuals choose b and the same occurs with type A individuals later on. On the right-hand side, slightly more than half of the population adopt. Most type B individuals eventually adopt the new technology, whereas most type A individuals choose the status quo. The reason that adoption is lower is that type A individuals are a greater fraction of the population in this case. This is consistent with our stability analysis in Section 3.

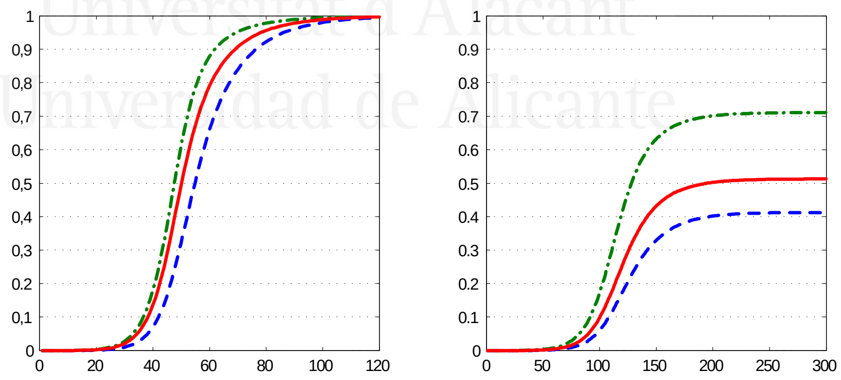

Figure 4: Adoption curves with full and less than full adoption

Less than full adoption and switching behavior. A distinctive feature of our model is that it allows for less than full adoption in equilibrium. Whenever this happens, there is a constant flow of individuals switching from the status quo to the new technology and vice versa. Indeed, what defines the interior equilibrium is the balance of these flows. At any state (p, q) the flow into the innovation is given by $f^+(t) := (\alpha(1-p(t))+(1-\alpha)q(t))(\alpha p(t)(1-L)+(1-\alpha)(1-q(t))R)$ and the flow away from it is $f^-(t) := (\alpha p(t) + (1-\alpha)(1-q(t)))(\alpha(1-p(t))L + (1-\alpha)q(t)(1-R))$ for all $t \geq 0$. Figure 5 provides three examples of equilibrium switching behavior. The parameters are set at $L = 0.6$, $R = 0.7$ and α equal to 0.558, 0.611 and 0.637. There are three sets of pairs of nearly parallel curves. The curves on the top of each pair represents $f^+(t)$ and the one on the bottom represents $f^-(t)$. Low final adoption corresponds to high values of α. The final fraction of adoption for each of these pair of curves is, going from bottom to top, 0.91, 0.71 and 0.61 and the corresponding fraction of individuals switching each time period is $0.078, 0.19$ and 0.218.[14] We can obtain the maximum flow to and from the innovation in equilibrium by maximizing $f^+(t)$, evaluated at $(p(t), q(t)) = (\hat{p}(L, R, \alpha), \hat{q}(R, L, 1-\alpha))$, with respect to L, R and α for all those triplets (L, R, α) such that $(\hat{p}(L, R, \alpha), \hat{q}(R, L, 1-\alpha))$ is an internal equilibrium. The numerical solution of this problem shows that these flows involve at most 25% of the population in equilibrium.[15]

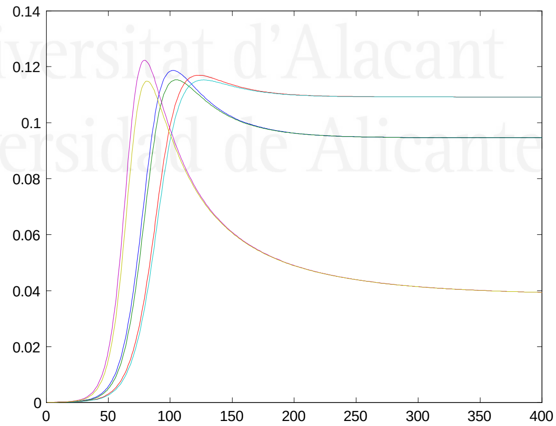

Switching in Equilibrium

[14] These numbers correspond to the sum of $f^+(t)$ and $f^-(t)$ once convergence is attained.

[15] The maximum is attained at $L = 0.500022$, $R = 0.500014$ and $\alpha = 0.384143$, or $R = 0.500022$, $L = 0.500014$ and $1 - \alpha = 0.384143$.

These results are consistent with a number of empirical findings in Suri's (2011) study of the adoption of hybrid corn in Kenya. He notes that aggregate adoption has flattened out at about 70% of the population since the 1990s, suggesting that adoption has reached an equilibrium level. Suri estimates production functions for agricultural output and found the return to adoption to be heterogenous across farmers. It is also found that some farmers switch back and forth between hybrid and non-hybrid corn. For example, between 1997 and 2004 around 39% of the households switched on and out in use. Suri suggests that the heterogeneity in adoption is due to the heterogeneity in returns and that close to indifferent farmers are responsible for the switching, since their decisions are sensitive to shocks to the costs of hybrid seed and fertilizers. Our model provides an alternative explanation for both facts. Here, less than full adoption is possible for a large set of parameter values. This occurs if there is a level of final adoption such that the flow of individuals switching into hybrid corn equals the flow of agents switching out. Hence, finding switching in equilibrium is a necessary consequence of less than full adoption in our model. Our model also provides an additional prediction. The fraction of adopters in the part of the population for which hybrid corn is not optimal, increases in the fraction of the population for which adoption is optimal. If for example infrastructure is improved, so hybrid corn becomes the optimal choice for a larger fraction of the population, then the fraction of adopters in the remaining part of the population should increase as well.

1.5 Biased Sampling

In this section we relax the assumption of uniform sampling and allow $\alpha_A \neq \alpha$ and $\alpha_B \neq 1 - \alpha$, so that each type may have stronger (or weaker) tendency to sample individuals of the same type. In other words, we allow for *homophily* (bias towards sampling same type individuals) and *heterophily* (bias towards sampling other type individuals) in the population. Homophilic tendencies are widely documented (see Currarini, Jackson and Pin 2009 and the references therein) and may be due to segregation or individual preferences for having friends that are similar to them. We introduce homophily and heterophily in our model allowing $\alpha_T = \frac{\varepsilon_T|T|}{\varepsilon_T|T|+|W\setminus T|}$ where $\varepsilon_T \in (0,\infty)$, is a constant that measures the bias of type T individuals to sample individuals of the same type for $T \in \{A, B\}$. If $\varepsilon_T = 1$ we obtain uniform sampling, whereas if $\varepsilon_T > 1$ type T individuals are homophilous, and if $\varepsilon_T < 1$, they are heterophilous.

Stable Equilibria. The resulting dynamics are similar to the case of uniform sampling, yet the analysis allows us to obtain insights about the role of these biases on the choices of the population. As before, we obtain convergence either to a corner rest point or to an internal rest point. Thus, the corresponding results that we obtained for uniform sampling are robust to biased sampling. However now we obtain some other results that cannot arise under uniform sampling. We show that convergence of the population to the non-optimal action for a type of individuals can always be ruled out if this type is sufficiently homophilous. More generally, the fraction of both types of individuals choosing a (correspondingly, b) increases in α_A (correspondingly α_B). This means that a type benefits from being more homophilous and is affected negatively by the homophily of the other type. In the limit, as α_A and α_B go to one, so each type is completely homophilous, the global attractor of the system goes to $(1,1)$, i.e., the entire population makes the right choice. Hence, the limit of the model corresponds to the case of two homogenous populations.

Formally, fix $L, R \in (\frac{1}{2}, \frac{3}{4}]$, let $\overline{\alpha}_A : \mathbb{R} \to \mathbb{R}$ such that $\overline{\alpha}_A(z) = \frac{R-L+zR(2L-1)}{(1-L)(2R-1)}$ for all $z \in \mathbb{R}$, and $\overline{\alpha}_B : \mathbb{R} \to \mathbb{R}$ such that $\overline{\alpha}_B(z) = \frac{L-R+zL(2R-1)}{(1-R)(2L-1)}$ for all $z \in \mathbb{R}$. It is easy to see that $\overline{\alpha}_A(\alpha_B) > \left[\overline{\alpha}_B^{-1}\right](\alpha_B)$ for all $\alpha_B \in (0,1)$, where $\left[\overline{\alpha}_B^{-1}\right]$ is the inverse function of $\overline{\alpha}_B$. The following lemma characterizes virtually all the pairs (α_A, α_B) such that $(1,0)$ is asymptotically stable and such that $(0,1)$ is asymptotically stable.[16]

Lemma 1.1. *Suppose $\alpha_A \neq \overline{\alpha}_A(\alpha_B)$ and $\alpha_B \neq \overline{\alpha}_B(\alpha_A)$, then (i) $(1,0)$ is asymptot-*

[16]The only possibilities we do not consider here are when $\alpha_A = \overline{\alpha}_A(\alpha_B)$ or $\alpha_B = \overline{\alpha}_B(\alpha_A)$, in which case some restpoints may not be hyperbolic and hence may not be determined using the Jacobian of the system.

ically stable if and only if $\alpha_A > \overline{\alpha}_A(\alpha_B)$, *(ii)* $(0,1)$ *is asymptotically stable if and only if* $\alpha_A < \left[\overline{\alpha}_B^{-1}\right](\alpha_B)$, *(iii) if* $(1,0)$ *is asymptotically stable, then* $(0,1)$ *is not asymptotically stable (and vice versa).*

Proof. (i) We use the determinant and trace of the Jacobian matrix. $Det(J(1,0)) = L - R + \alpha_B R(1-2L) + \alpha_A(2R-1)(1-L)) > 0$ iff $\alpha_A > \overline{\alpha}_A(\alpha_B)$. $\overline{\alpha}_A(\alpha_B) \geq 1$ for all $\alpha_B \geq \frac{1-R}{R}$, so if $\alpha_A > \overline{\alpha}_A(\alpha_B)$, then $\alpha_B < \frac{1-R}{R}$. Finally, if $\alpha_B < \frac{1-R}{R}$, then $Tr(J(1,0)) = \alpha_A - 1 - L(1+\alpha_A) + R(1+\alpha_B) < 0$. (ii) is established analogously. (iii) follows from (i) and (ii) and the fact that $\overline{\alpha}_A(\alpha_B) > \left[\overline{\alpha}_B^{-1}\right](\alpha_B)$ for all $\alpha_B \in (0,1)$. □

Corollary 3. *(i) If* $\alpha_B > \frac{1-R}{R}$ *then* $(1,0)$ *is not asymptotically stable, and if* $L > R$ *and* $\alpha_B < \frac{L-R}{R(2L-1)}$, *then* $(1,0)$ *is asymptotically stable. (ii) If* $\alpha_A > \frac{1-L}{L}$ *then* $(0,1)$ *is not asymptotically stable, and if* $R > L$ *and* $\alpha_A < \frac{R-L}{L(2R-1)}$, *then* $(0,1)$ *is asymptotically stable.*

Part (i) of Corollary 3 follows by observing that if $\alpha_B > \frac{1-R}{R}$ then $\overline{\alpha}_A(\alpha_B) > 1$ and if $L > R$ and $\alpha_B < \frac{L-R}{R(2L-1)}$, then $\overline{\alpha}_A(\alpha_B) < 0$ (and the analogous argument proves (ii)).

Lemma 1.1 implies that for any $L, R \in (\frac{1}{2}, \frac{3}{4}]$ and $\alpha_A, \alpha_B \in (0,1)$ either (i) $(1,0)$ is asymptotically stable, (ii) $(0,1)$ is asymptotically stable, or (iii) neither is asymptotically stable. For large values of α_A relative to α_B, $(1,0)$ is asymptotically stable, whereas for large values of α_B relative to α_A, $(0,1)$ is asymptotically stable. For more similar values of α_A and α_B, neither $(1,0)$, nor $(0,1)$ is asymptotically stable. This is illustrated in Figure 6 for $L = 0.7$ and $R = 0.6$.

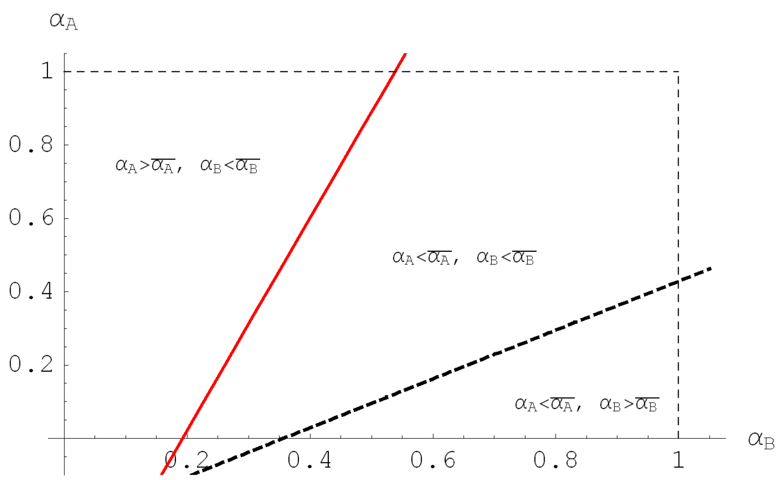

Figure 6: The solid line corresponds to $\overline{\alpha}_A$ and the dashed line to $\overline{\alpha}_B$.

Furthermore, if and only if $\left[\overline{\alpha}_B^{-1}\right](\alpha_B) < \alpha_A < \overline{\alpha}_A(\alpha_B)$ (i.e., if neither $(1,0)$ nor $(0,1)$ are asymptotically stable) there is also a unique internal rest point (for details, see Appendix D). The following results show that the asymptotically stable rest points are global attractors.

Proposition 1.5. *If* $(p(0), q(0)) \notin RP$, *then* $\lim_{t\to\infty}(p(t), q(t)) = (p^*, q^*)$, *where*

$$(p^*, q^*) \begin{cases} = (0,1) & \textit{if} \quad \alpha_A < \left[\overline{\alpha}_B^{-1}\right](\alpha_B) \\ \in (0,1)^2 & \textit{if} \quad \left[\overline{\alpha}_B^{-1}\right](\alpha_B) < \alpha_A < \overline{\alpha}_A(\alpha_B) \\ = (1,0) & \textit{if} \quad \overline{\alpha}_A(\alpha_B) < \alpha_A. \end{cases}$$

Proof. Since $\dot{p}_2(p,q), \dot{q}_2(q,p) < 0$ the system always converges to a rest point as $t \to \infty$. If $\alpha_A > \overline{\alpha}_A$ then by Remark 11 and Lemma 1.11, $RP = \{(1,0),(0,1)\}$. By Lemma 1.1, $(0,1)$ is unstable and has no stable arm in $[0,1]^2$. Hence, the system converges to $(1,0)$. The converse holds for $\alpha_B > \overline{\alpha}_B$. If $\alpha_A < \overline{\alpha}_A$ and $\alpha_B < \overline{\alpha}_B$, then by Lemma 1.11, $RP = \{(1,0),(0,1),(\hat{p},\hat{q})\}$, with $(\hat{p},\hat{q}) \in (0,1)^2$. By Lemma 1.1, $(1,0)$ and $(0,1)$ are unstable and have no stable arm in $(0,1)^2$. Hence, the system converges to $(\hat{p},\hat{q})$. □

In Figure 6, the northwest (correspondingly southeast) region corresponds to the parameter values such that $(1, 0)$ (correspondingly $(0, 1)$) is a global attractor. In the center region the internal rest point is the global attractor.

We shall emphasize two implications of the results above. First, together with Corollary 3, Proposition 1.5 implies that if individuals of a given type are sufficiently homophilous, then the population will not converge to choose the action that is not optimal for that type of individuals. For example, if $\alpha_B > \frac{1-R}{R}$ then the system will not converge to $(1, 0)$ *regardless of* α_A and L. An implication is that if both types are sufficiently homophilous, then the population always converge to the internal rest point. Second, if the bandit of type B individuals is less risk sensitive than the bandit of type A individuals, and additionally type B is relatively small or heterophilous, so α_B is below a threshold value (determined by L and R), then the population converges to choose a, *regardless of* α_A. This is seen clearly in Figure 6, where if α_B is below around 0.15, then the population converges to choose a even if α_A is arbitrarily small.

Comparative Statics. The more likely is a type of individuals to sample individuals of their own type, the larger the fraction of that type that chooses its optimal action in an internal rest point. Fix $L, R \in \left(\frac{1}{2}, \frac{3}{4}\right]$ and let $(\hat{p}(\alpha_A, \alpha_B), \hat{q}(\alpha_B, \alpha_A))$ be the internal rest point, when it exists, as a function of α_A and α_B. Using implicit differentiation, it is easy to establish that $\hat{p}_1(\alpha_A, \alpha_B) > 0$, $\hat{p}_2(\alpha_A, \alpha_B) < 0$, $\hat{q}_1(\alpha_B, \alpha_A) > 0$ and $\hat{q}_2(\alpha_B, \alpha_A) < 0$ (see Appendix D). The intuition of this result follows from the fact that internal rest points are characterized by the fact that flows in to and out of each action for each type are balanced. Since internal rest points (p^*, q^*) always satisfy $(p^*, q^*) \gg (\frac{1}{2}, \frac{1}{2})$, in equilibrium, a higher tendency to sample own-type individuals, makes it more likely to sample an individual playing the action that is optimal for both types. This allows that in equilibrium, to maintain the balance, there are a greater fraction of the population playing the already the optimal action (and hence likely to be an outflow) and a smaller fraction playing the non-optimal action (and hence likely to be an inflow).

Furthermore, within the internal rest points, as the probability of a type of sampling individuals of their own type goes to one, in equilibrium, the fraction of those individuals who choose their optimal action goes to one. This is formalized in the following remark (the proof is provided in Appendix D).

Remark 4. *If* $\alpha_B > \frac{1-R}{R}$, *then* $\lim_{\alpha_A \to 1} \hat{p}(\alpha_A, \alpha_B) = 1$ *and if* $\alpha_A > \frac{1-L}{L}$, *then* $\lim_{\alpha_B \to 1} \hat{q}(\alpha_B, \alpha_A) = 1$.

The role of the qualifiers in the remark is just to guarantee that we are in a case where there is an internal rest point. Finally, notice that this result implies that a heterogenous population behaves as two homogenous population, in the limit, as each type becomes completely homophilous.

1.6 Discussion

The model of social comparison and learning we analyze in this paper allows us to obtain sharp results. In a two-type population where each type faces a bandit with relatively similar risk sensitivities, the population is lead towards the action that is optimal for the majority but non-optimal for the minority. This result is obtained even though individuals of different types are somehow aware of the fact that different actions may lead to different payoffs for different individuals. This is less likely to occur if the population is relatively balanced. In such a case, the population may converge to an asymptotically stable state in which most individuals of both types choose the action that is optimal for each type.

Despite the fact that a type of individuals may be lead to their non-optimal action in the steady state, the average expected payoff of the population in every asymptotically stable point is greater than the expected payoff of a heterogeneous population of individuals who choose their actions randomly. Intuitively, when the type of individuals converging to its non-optimal action is relatively small or faces a bandit that is not very risk-sensitive, the losses associated with not choosing the optimal action are relatively small.

Overall, the model of social learning and comparison can suitably be applied to diffusion of innovations when the returns to adoption are heterogenous. The rates of adoption, L and R, reflect the relative merit of the innovation and the status quo for the corresponding types in the population. The adoption curves tend to have the typical S-form. Final outcomes of less than full adoption are fully consistent with this process of diffusion. Equilibrium switching behavior arises as a natural consequence of any such outcome.

In our analysis of a heterogeneous population we assume that individuals are equally likely to observe the choice and payoff of any other individual. Indeed, depending on the degree of segregation of a population, individuals may be more or less likely to observe the choices and outcomes of individuals that are different from them. Our results suggest that minorities are better off when there is some segregation, than in integrated heterogeneous population.

There are several extensions of the model we have left for future research. In our setup there is no role for the accuracy of the signal. It is intuitive that the experiences of people who are perceived as different may be less informative, and indeed, some empirical evidence suggests that information about different individuals is often discarded (e.g., Munshi, 2004). In our model, however, the information about different individuals is as informative as the information about similar individuals, the only difference is that, since the comparison-signal is random, yet unbiased, the perceived payoff from the sampled individual (the payoff he obtained plus the comparison-signal) is a mean preserving spread of the perceived payoff from the sampled individual if the signal were not assumed to be noisy. Nevertheless, as the expected probability of playing each action is not affected by mean preserving spreads, in terms of the expected probability of choosing each action, the randomness of the comparison-signal does not play an important role in our analysis.

Our model assumes as given the exogenous sampling process. However, in the presence of heterogeneity individuals may have incentives to search for individuals that are similar to them and therefore suitable to learn from. At the same time, it seems that individuals would prefer to sample others who have made good choices. An analysis that considers an endogenous sampling process might provide interesting insights into the implications of heterogeneity in the search for suitable role model.

We leave for future research a more thorough consideration of these issues.

1.7 Appendix A: Proof of Necessity in Proposition 1.1

Necessity in Proposition 1.1 is argued using the following lemmata.

Lemma 1.2. *If L is payoff-ordering, then, if $\pi_{\tau d} = \pi_{\tau c}$, then $L^d_{cd}(\tau, \tau') = 1/2$.*

Proof. Consider an environment F such that $\pi_{\tau d} = \pi_{\tau c}$ and assume $L^d_{cd}(\tau, \tau') < 1/2$. We will now consider a small ε perturbation of this environment, which we denote $\widetilde{F}$. First we perturb $F_{\tau d}$ by letting $\widetilde{F}_{\tau d}$ be such that for any interval $I \subset [0, 1)$ we have $\widetilde{\mu}_{\tau d}(I) = (1 - \varepsilon)\mu_{\tau d}(I)$ and $\widetilde{\mu}_{\tau d}(1) = \mu_{\tau d}(1) + \varepsilon\mu_{\tau d}[0, 1)$. We obtain $\widetilde{\pi}_{\tau d} = (1 - \varepsilon)\pi_{\tau d} + \varepsilon$. Next, we perturb $F_{\tau c}$ by letting $\widetilde{F}_{\tau c}$ be such that for any $I \subset (0, 1]$ we have $\widetilde{\mu}_{\tau c}(I) = (1 - \varepsilon)\mu_{\tau c}(I)$ and $\widetilde{\mu}_{\tau c}(0) = \mu_{\tau c}(0) + \varepsilon\mu_{\tau c}(0, 1]$. We obtain $\widetilde{\pi}_{\tau c} = (1 - \varepsilon)\pi_{\tau c}$ and

$$
\begin{aligned}
E[\widetilde{\delta}_{\tau\tau' d}] &= (1 - \varepsilon)\pi_{\tau d} + \varepsilon - \pi_{\tau d} \\
&= (1 - \varepsilon)\left(\pi_{\tau d} - \pi_{\tau' d}\right) + \varepsilon(1 - \pi_{\tau' d}) \\
&= (1 - \varepsilon)E\left[\delta_{\tau\tau' d}\right] + \varepsilon(1 - \pi_{\tau' d}).
\end{aligned}
$$

Suppose the distribution of $\widetilde{\delta}_{\tau\tau' d}$ is given by a compounded distribution which weights with probabilities $1 - \varepsilon$ and ε the distributions $F_{\delta_{\tau\tau' d}}$ and a degenerate distribution which assigns all the probability mass to $1 - \pi_{\tau' d}$. In all the other respects, the environments F and $\widetilde{F}$ are the same. Let $\widetilde{L}^d_{cd}(\tau, \tau')$ denote the expected updated probability of choosing d when $i \in \tau$ chooses c, observes $j \in \tau'$ who chooses d, and the comparison-signal $\widetilde{\delta}_{\tau\tau' d}$ in the environment $\widetilde{F}$. $\widetilde{L}^d_{cd}(\tau, \tau')$ can be written as a continuous function of ε and when $\varepsilon = 0$, $\widetilde{L}^d_{cd}(\tau, \tau') = L^d_{cd}(\tau, \tau') < 1/2$. Thus, for small enough ε, $\widetilde{L}^d_{cd}(\tau, \tau') < 1/2$ and, since $\widetilde{\pi}_{\tau d} > \widetilde{\pi}_{\tau c}$, L is not payoff-ordering. □

Corollary 5. *If L is payoff-ordering, $x, y \in [1, -1]$, $\delta \in [-1, 1]$, and $x = y + \delta$ then, $L(c, x, d, y, \delta)(d) = 1/2$.*

Proof. Consider x, y, δ which satisfy the hypothesis and an environment in which $\mu_{\tau c}(x) = \mu_{\tau d}(y+\delta) = \mu_{\tau' d}(y) = \mu_{\delta_{\tau\tau' d}}(\delta) = 1$. In this environment $L^d_{cd}(\tau, \tau') = L(c, x, d, y, \delta)(d)$, and thus, the previous lemma implies $L(c, x, d, y, \delta)(d) = 1/2$. □

Lemma 1.3. *If L is payoff-ordering, then,*

$$L(c, x-\varepsilon, d, y, \delta)(d) = L(c, x, d, y+\varepsilon, \delta)(d) = L(c, x, d, y, \delta+\varepsilon)(d)$$

for all $x, y \in [1, -1]$, $\delta \in [-1, 1]$, and ε such that $x-\varepsilon, y+\varepsilon \in [0, 1]$ and $\delta+\varepsilon \in [-1, 1]$.

Proof. Consider an environment in which $\mu_{\tau c}(x) = \mu_{\tau c}(x-\varepsilon) = 1/2$, $\mu_{\tau d}(x-\varepsilon/2) = \mu_{\tau' d}(y) = 1$ for some x, y, ε such that $x-\varepsilon, y+\varepsilon \in [0, 1]$ and $\delta + y + \varepsilon/2 - x \neq 0$. Let $\mu_{\delta_{\tau\tau' d}}(\delta) = p$ and $\mu_{\delta_{\tau\tau' d}}(\delta') = 1-p$, for some $\delta > x-y-\varepsilon/2$ and $\delta' < x-y-\varepsilon/2$ such that $\delta+\varepsilon, \delta'+\varepsilon \in [-1, 1]$ and $p\delta + (1-p)\delta' = \pi_{\tau d} - \pi_{\tau' d} = x - y - \varepsilon/2$, which implies $p = \frac{x-y-\delta'-\varepsilon/2}{\delta-\delta'}$ and $1-p = \frac{\delta+y+\varepsilon/2-x}{\delta-\delta'}$. Payoff ordering imposes $pL(c, x, d, y, \delta)(d) + pL(c, x-\varepsilon, d, y, \delta)(d) + (1-p)L(c, x, d, y, \delta')(d) + (1-p)L(c, x-\varepsilon, d, y, \delta')(d) = 1$. Now consider an environment in which $\mu_{\tau' d}(y) = \mu_{\tau' d}(y+\varepsilon) = 1/2$, $\mu_{\tau c}(x) = \mu_{\tau d}(x) = 1$, $\mu_{\delta_{\tau\tau' d}}(\delta) = p$ and $\mu_{\delta_{\tau\tau' d}}(\delta') = 1-p$. Payoff ordering imposes $pL(c, x, d, y, \delta)(d) + pL(c, x, d, y+\varepsilon, \delta)(d) + (1-p)L(c, x, d, y, \delta')(d) + (1-p)L(c, x, d, y+\varepsilon, \delta')(d) = 1$. Subtract this from the equality in the preceding paragraph to obtain $p\frac{L(c,x-\varepsilon,d,y,\delta)(d)-L(c,x,d,y+\varepsilon,\delta)(d)}{p-1} = L(c, x-\varepsilon, d, y, \delta')(d) - L(c, x, d, y+\varepsilon, \delta')(d)$ and $\frac{L(c,x-\varepsilon,d,y,\delta)(d)-L(c,x,d,y+\varepsilon,\delta)(d)}{\delta+y+\varepsilon/2-x} = \frac{L(c,x-\varepsilon,d,y,\delta')_d-L(c,x,d,y+\varepsilon,\delta')(d)}{\delta'+y+\varepsilon/2-x}$. Note that the left and right hand side of this expression are independent of δ and δ'. By the preceding Corollary, if we take $\delta = x-y-\varepsilon$ (which satisfies $\delta+\varepsilon = x-y \in [-1, 1]$) we obtain $L(c, x-\varepsilon, d, y, x-y-\varepsilon)(d) = L(c, x, d, y+\varepsilon, x-y-\varepsilon)(d) = 1/2$, and hence, $L(c, x-\varepsilon, d, y, \delta)(d) - L(c, x, d, y+\varepsilon, \delta)(d) = 0$.

Now, let $\mu_{\tau c}(x) = \mu_{\tau d}(x) = \mu_{\tau' d}(y) = 1$. Let $\mu_{\delta_{\tau\tau' d}}(\delta) = p/2$, $\mu_{\delta_{\tau\tau' d}}(\delta+\varepsilon) = p/2$ and $\mu_{\delta_{\tau\tau' d}}(\delta') = (1-p)/2$ and $\mu_{\delta_{\tau\tau' d}}(\delta'+\varepsilon) = (1-p)/2$. Then, $E[\delta_{\tau\tau' d}] = p\delta + p\varepsilon/2 + (1-p)\delta' + (1-p)\varepsilon/2 = p\delta + (1-p)\delta' + \varepsilon/2 = x-y$ and $pL(c, x, d, y, \delta)(d) + pL(c, x, d, y, \delta+\varepsilon)(d) + (1-p)L(c, x, d, y, \delta')(d) + (1-p)L(c, x, d, y, \delta'+\varepsilon)(d) = 1$. Combining this with the expression above we obtain $p(L(c, x-\varepsilon, d, y, \delta)(d) - L(c, x, d, y, \delta+\varepsilon))(d) + (1-p)(L(c, x-\varepsilon, d, y, \delta')(d) - L(c, x, d, y, \delta'+\varepsilon)(d)) = 0$, or $\frac{L(c,x-\varepsilon,d,y,\delta)(d)-L(c,x,d,y,\delta+\varepsilon)(d)}{(\delta+y+\varepsilon/2-x)} = \frac{L(c,x-\varepsilon,d,y,\delta')(d)-L(c,x,d,y,\delta'+\varepsilon)(d)}{(\delta'+y+\varepsilon/2-x)}$. Again, take $\delta = x-\varepsilon-y$, which yields $L(c,x-\varepsilon,d,y,x-\varepsilon-y)(d) = L(c,x,d,y,x-y)(d) = 1/2$, and therefore, $L(c,x-\varepsilon,d,y,\delta)(d) - L(c,x,d,y,\delta+\varepsilon)(d) = 0$.

The results for $\delta = y+\varepsilon/2-x$ are obtained by considering an environment in which $\delta = \delta'$ and doing analogous computations. □

Corollary 6. *For all $x, y, x', y' \in [0,1]$, and $\delta, \delta' \in [-1,1]$ such that $y+\delta-x = y'+\delta'-x'$, we have $L(c,x,d,y,\delta)(d) = L(c,x',d,y',\delta')(d)$.*

Proof. Note that in this case

$$\begin{aligned} L(c,x,d,y,\delta)(d) &= L(c,x'+y-y'+\delta-\delta',d,y,\delta)(d) \\ &= L(c,x',d,y-y+y',\delta-\delta+\delta')(d) \\ &= L(c,x',d,y',\delta')(d), \end{aligned}$$

where the second equality follows from the preceding lemma. □

This result implies that $L(c,x,d,y,\delta)(d)$ can be written as a function of $y+\delta-x$ only. Let this function be denoted by $g : [-2,2] \to [0,1]$. Corollary 1 implies $g(0) = 1/2$. The following results show that g is also linear, i.e., $g(z) = 1/2 + \beta z$ for some real number β and all $z \in [-2,2]$.

Lemma 1.4. *If L is payoff-ordering then, $g(-z) + g(z) = 1$ for $z \in [-1,1]$.*

Proof. Let $\mu_{\tau c}(0) = \mu_{\tau d}(0) = \mu_{\tau' d}(0) = 1$ and $\mu_{\delta_{\tau\tau' d}}(-z) = \mu_{\delta_{\tau\tau' d}}(z) = 1/2$, $z \in [0,1]$. Then, payoff ordering imposes $1/2g(-z) + 1/2g(z) = 1/2$. □

Lemma 1.5. *If L is payoff-ordering then $g(-z) + g(z) = 1$ for $z \in [-2,2]$.*

Proof. Let $\mu_{\tau c}(x) = \mu_{\tau c}(0) = \mu_{\tau' d}(x) = \mu_{\tau' d}(0) = 1/2$, $\mu_{\tau d}(x/2) = 1$ and $\mu_{\delta_{\tau\tau' d}}(-x) = \mu_{\delta_{\tau\tau' d}}(x) = 1/2$, $x \in [0,1]$. Payoff ordering imposes

$$\frac{1}{8}g(-2x) + \frac{2}{8}g(-x) + \frac{2}{8}g(0) + \frac{2}{8}g(x) + \frac{1}{8}g(2x) = 1/2.$$

This expression and the preceding lemma imply $g(-2x) + g(2x) = 1$, which along with the preceding lemma yield the result. □

Corollary 7. *If L is payoff-ordering then $g(0) = 1/2g(-z) + 1/2g(z)$.*

Lemma 1.6. *If L is payoff-ordering then $g(z) = 1/2g(z+\varepsilon) + 1/2g(z-\varepsilon)$ for $z \in [-1,1]$ and $\varepsilon \in [0, 1-|z|]$.*

Proof. Let $\mu_{\tau c}(0) = \mu_{\tau d}(0) = \mu_{\tau' d}(0) = 1$ and $\mu_{\delta_{\tau\tau' d}}(-z) = 1/2$ and $\mu_{\delta_{\tau\tau' d}}(z+\varepsilon) = \mu_{\delta_{\tau\tau' d}}(z-\varepsilon) = 1/4\,,\ z \in [0,1]$ and $\varepsilon \in [0, 1-z]$. Then, payoff ordering implies

$$
\begin{aligned}
\frac{1}{2}g(-z) + \frac{1}{4}g(z-\varepsilon) + \frac{1}{4}g(z+\varepsilon) &= \frac{1}{2} \\
g(-z) + \frac{1}{2}(g(z+\varepsilon) + \frac{1}{2}g(z-\varepsilon)) &= 1.
\end{aligned}
$$

By substituting this in $g(-z) + g(z) = 1$ we obtain $g(z) = 1/2g(z+\varepsilon) + 1/2g(z-\varepsilon)$. □

A consequence of the preceding lemma is that payoff ordering requires g to be linear over the interval $[-1,1]$. Linearity follows from the standard result that if $g(1/2z + 1/2z') = 1/2g(z) + 1/2g(z')$ for all z, z' in some interval $[\underline{z}, \overline{z}]$ then, g is both concave and convex over $[\underline{z}, \overline{z}]$. In the following lemma we show that linearity holds over the whole domain of g.

Lemma 1.7. *If L is payoff-ordering then, $g(z) = 1/2g(z+\varepsilon) + 1/2g(z-\varepsilon)$ for $z \in [-2,1] \cup [1,2]$ and $\varepsilon \in [0, 1-|z|]$.*

Proof. Let $\mu_{\tau c}(0) = \mu_{\tau' d}(0) = \mu_{\tau' d}(x) = 1/2$, $\mu_{\tau c}(x+\varepsilon) = \mu_{\tau c}(x-\varepsilon) = 1/4$, $\mu_{\delta_{\tau\tau' d}}(-1) = \mu_{\delta_{\tau\tau' d}}(1) = 1/2$ and $\mu_{\tau d}(\pi_{\tau c}) = 1$, with $x \in [0,1]$ and $\varepsilon \in [0, 1-x]$. Then, by using payoff ordering, symmetry around zero and by applying the preceding lemma, it can be derived that $g(x+1) = 1/2g(x+1+\varepsilon) + 1/2g(x+1-\varepsilon)$ which implies $g(z) = 1/2g(z+\varepsilon) + 1/2g(z-\varepsilon)$ for $z \in [-2,1] \cup [1,2]$ and $\varepsilon \in [0, 1-|z|]$. □

This lemma implies that $g(z)$ is linear in $[-2,1] \cup [1,2]$. Hence, the two preceding lemmata imply linearity in $[-2,2]$. Consequently, a payoff-ordering decision rule must satisfy $L(c,x,d,y,\delta)_d = 1/2 + \beta(y+\delta-x)$. Consider $x, y \in [0,1]$ and $\delta \in [-1,1]$ such that $y+\delta > x$ and the environment F such that $\mu_{\tau c}(x) = \mu_{\tau d}(y+\delta) = \mu_{\tau' d}(y) = \mu_{\delta_{\tau\tau' d}}(\delta) = 1$. Since $L^d_{cd} = L(c,x,d,y,\delta)_d = 1/2 + \beta(y+\delta-x)$, payoff ordering implies that $\beta > 0$. Finally, since the range of L is $[0,1]$ and in the domain $x, y \in [0,1]$ and $\delta \in [-1,1]$, we also need $\beta \leq 1/4$.

1.8 Appendix B: Properties of Payoff-Ordering Rules

Claim 8. *For any decision rule L and profile of choices $(c(i))_{i\in W}$, there is an environment F such that*

$$\int\int p_i(j) L^{c(j)}_{c(i)c(j)}(\tau(i),\tau(j))\pi_{\tau(i)c(j)}\,dj\,di < \int \pi_{\tau(i)c(i)}\,di$$

The (omitted) proof follows from the simple observation that for every (non-trivial) improving rule in an homogeneous population such that $(c(i))_{i\in W}$ satisfies $c(i) \in \arg\max_{c\in S}\{\pi_{\tau(i)c}\}$ for all $i \in W$, there is an environment in which, with strictly positive probability, some individual i observes another individual j and changes her action from $c(i)$ to $c(j)$ although $c(j) \notin \arg\max_{c\in S}\{\pi_{\tau(i)c}\}$.

Claim 9. *If $\pi_{ic} = \pi_{jc}$ for all $c \in S$ and $p_{ij} = p_{ji}$, for all $(i,j) \in W^2$, then for any payoff-ordering decision rule L and profile of decisions $(c(i))_{i\in W}$,*

$$\int\int p_i(j) L^{c(j)}_{c(i)c(j)}(\tau(i),\tau(j))\pi_{\tau(i)c(j)}\,dj\,di \geq \int \pi_{\tau(i)c(i)}\,di.$$

The omitted argument is quite intuitive and it follows from the fact that every individual who may change her action is more likely to change to a an action with a greater payoff (when, as it is usually assumed in the analysis of homogeneous populations, matching probabilities are symmetric).

1.9 Appendix C: Proof of Proposition 1.3

Proof. First we establish asymptotic stability. Let

$$J(p,q) := \begin{bmatrix} \dot{p}_1(p,q) & \dot{p}_2(p,q) \\ \dot{q}_2(q,p) & \dot{q}_1(q,p) \end{bmatrix},$$

where $\dot{p}_i$ and $\dot{q}_i$ denote their corresponding partial derivatives with respect to their i^{th} argument. A rest point (p^*, q^*) is asymptotically stable if the real part of the eigenvalues of $J(p^*, q^*)$ are negative. This is equivalent to $Det(J(p^*, q^*)) > 0$ and $Tr(J(p^*, q^*)) < 0$, where $Det(J(p^*, q^*))$ and $Tr(J(p^*, q^*))$ are the determinant and trace of $J(p^*, q^*)$ respectively. Consider first $(1, 0)$. We have $Tr(J(1, 0)) = 2R-(1+L)-\alpha(L+R-1) < 0$. Next, $Det(J(1, 0)) = L(1-2R)+\alpha(-1+L+R) > 0$ is equivalent to $\alpha > \frac{L(1-2R)}{L+R-1} = \overline{\alpha}(L, R)$. An analogous calculation holds for $(0, 1)$. Now, consider $(\hat{p}(L, R, \alpha), \hat{q}(R, L, 1-\alpha))$. Note that $Tr(J(\hat{p}(L, R, \alpha), \hat{q}(R, L, 1-\alpha))) = \frac{2L(1-L)(\alpha+(1-2R)^2)+2R(1-R)(2-\alpha)-1}{(2R-1)(2L-1)} < 0$. Next, $Det(J(\hat{p}(L, R, \alpha), \hat{q}(R, L, 1-\alpha))) = \frac{(\alpha(L+R-1)-(1-L)(2R-1))(L(2R-1)-\alpha(L+R-1))}{(2L-1)(2R-1)} > 0$ if $\alpha > \frac{(1-L)(2R-1)}{(L+R-1)} = \underline{\alpha}(L, R)$ and $\alpha < \frac{L(2R-1)}{(L+R-1)} = \overline{\alpha}(L, R)$.

In order to prove that the asymptotically stable points are global attractors, notice that $\dot{p}_2, \dot{q}_2 < 0$, hence $\lim_{t\to\infty}(p(t), q(t)) \in RP$ for all paths. Finally for all $(p, q) \in RP$ which are not asymptotically stable, both eigen values of $J(p, q)$ are positive. Hence there is no neighborhood around (p, q) that contains a path converging to it. □

1.10 Appendix D: Phase Diagram Analysis

In this section we derive the phase diagrams corresponding to the system (1.3)-(1.4). This is done by using equations (1.5)-(1.6) and observing that $\dot{p} \lesseqqgtr 0 \iff q \gtreqqless q(p)$ and $\dot{q} \lesseqqgtr 0 \iff p \gtreqqless p(q)$. In other words, the direction in which the system (1.3)-(1.4) moves in (p, q) space given some initial point $(p(0), q(0))$ will depend on whether this point is above or below the graphs of the functions $\overline{p}(\cdot)$ and $\overline{q}(\cdot)$ defined in the proof of Proposition 1.2. In what follows we will characterize the behavior of the dynamics in three different regions of $[0, 1]^2$, which allows us to make conclusions about the long run convergence of the process. Define the function $[p^{-1}] : [0, 1] \to [0, 1]$ with $[p^{-1}](p) = q$ where $\overline{p}(q) = p$, for $p \in [0, 1)$; and $[p^{-1}](1) = \max\{q \in [0, 1] : \overline{p}(q) = 1\}$.[17] $[p^{-1}]$ is continuous, decreasing and concave. The following lemma is obtained by straightforward differentiation.

Lemma 1.8. *$\alpha < \underline{\alpha}(L, R)$ if and only if $\overline{q}'(0) < [p^{-1}]'(0)$.*

[17] Notice that $\{q : \overline{p}(q) = p\}$ is a singleon for $p \in [0, 1)$.

Lemma 1.9. *(i) If $\alpha < \underline{\alpha}(L,R)$, then for all $p \in (0,1)$ we have $\overline{q}(p) < [p^{-1}](p)$.*

(ii) If $\alpha > \bar{\alpha}(L,R)$, then for all $p \in (0,1)$ we have $\overline{q}(p) > [p^{-1}](p)$.

(iii) If $\alpha \in (\underline{\alpha}(L,R), \bar{\alpha}(L,R))$, then for $p \in (0,p^*)$ we have $\overline{q}(p) < [p^{-1}](p)$ and for $p \in (p^*,1)$ we have $\overline{q}(p) < [p^{-1}](p)$.

Proof. (i) Since $\alpha < \underline{\alpha}(L,R)$ we have $\overline{q}'(0) < [p^{-1}]'(0)$. Given that $[p^{-1}](0) = \overline{q}(0) = 1$, this means that for some (small) p we have $\overline{q}(p) < [p^{-1}](p)$. Since with $\alpha < \bar{\alpha}(L,R)$ there is no point $p \in (0,1)$ such that $\overline{q}(p) = [p^{-1}](p)$, the continuity of $\overline{q}(p)$ and $[p^{-1}]$ implies $\overline{q}(p) < [p^{-1}](p)$ for all $p \in (0,1)$. (ii) is established by an analogous argument. (iii) Since $\alpha \in (\underline{\alpha}(L,R), \bar{\alpha}(L,R))$ we have $\overline{q}'(0) > [p^{-1}]'(0)$. This means that for some (small) p we have $\overline{q}(p) > [p^{-1}](p)$. Given that there is a single point of intersection $(p^*,q^*) \in (0,1)$ at which $\overline{q}(p^*) = [p^{-1}](p^*)$, it must hold that $\overline{q}(p) > [p^{-1}](p)$ for all $p \in (0,p^*)$. By a similar argument it is established that $\overline{q}(p) < [p^{-1}](p)$ for all $p \in (p^*,1)$. □

Since $[p^{-1}](p)$ is decreasing if for (p,q) we have $q < [p^{-1}](p)$ then $\overline{p}(q) > p$. This implies that for (p,q) such that $q \lesseqgtr [p^{-1}](p)$ we have $\dot{q} \gtreqless 0$. Using this argument the following can be established.

Corollary 10. *(i) If $\alpha < \underline{\alpha}(L,R)$, then for any $p \in (0,1)$ and $q \in (\overline{q}(p), [p^{-1}](p))$ we have $\dot{p} < 0$ and $\dot{q} > 0$.*

(ii) If $\alpha > \bar{\alpha}(L,R)$, then for any $p \in (0,1)$ and $q \in ([p^{-1}](p), \overline{q}(p))$ we have $\dot{p} > 0$ and $\dot{q} < 0$.

(iii) If $\alpha \in (\underline{\alpha}(L,R), \bar{\alpha}(L,R))$, then for $p \in (0,p^*)$ and $q \in ([p^{-1}](p), \overline{q}(p))$ we have $\dot{p} > 0$ and $\dot{q} < 0$ and for $p \in (p^*,1)$ and $q \in (\overline{q}(p), [p^{-1}](p))$ we have $\dot{p} < 0$ and $\dot{q} > 0$.

Using these results we partition $[0,1]^2$ in four subsets. First, let $\mathcal{F}_1 := \{(p,q) \in [0,1]^2 : \dot{p}, \dot{q} > 0\}$, $\mathcal{F}_2 := \{(p,q) \in [0,1]^2 : \dot{p}, \dot{q} < 0\}$. Notice that $\mathcal{F}_1 = \{(p,q) \in [0,1]^2 : q < \min\{\overline{q}(p), p^{-1}(p)\}\}$ and $\mathcal{F}_2 = \{(p,q) \in [0,1]^2 : q > \max\{\overline{q}(p), p^{-1}(p)\}\}$,i.e. they are located to the southwest and northeast ,respectively, in the $[0,1]^2$ plane, and they cannot be empty since $(0,0) \in \mathcal{F}_1$ and $(1,1) \in \mathcal{F}_2$. Let $\mathcal{F}_3 := \{(p,q) \in [0,1]^2 :$

$q \in (\min\{q(p), p^{-1}(p)\}, \max\{q(p), p^{-1}(p)\})\}$, i.e., the area between $q(p)$ and $p^{-1}(p)$. The lemma implies that $\mathcal{F}_3$ is non-empty. We also know that if $\alpha < \underline{\alpha}(L, R)$, then at any $(p, q) \in \mathcal{F}_3$ it holds that $\dot{p} < 0$ and $\dot{q} > 0$. If $\alpha > \bar{\alpha}(L, R)$ then for any $(p, q) \in \mathcal{F}_3$ we have $\dot{p} < 0$ and $\dot{q} > 0$. Finally, if $\alpha \in (\underline{\alpha}(L, R), \bar{\alpha}(L, R))$ then for any $(p, q) \in \mathcal{F}_3$ such that $p \in (0, p^*)$ we have $\dot{p} > 0$ and $\dot{q} < 0$ while at $p \in (p^*, 1)$ we have $\dot{p} < 0$ and $\dot{q} > 0$. To complete the partition, let $\mathcal{F}_4 := \{(p, q) \in [0, 1]^2 : q = q(p) \vee p = [p^{-1}](p)\}$, and notice that $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3 \cup \mathcal{F}_4 = [0, 1]^2$. $\mathcal{F}_3 \cup \mathcal{F}_4$ "separates" $\mathcal{F}_1$ from $\mathcal{F}_2$, in the sense that for any continuous path $(p(t), q(t))$ in which $(p(t'), q(t')) \in \mathcal{F}_1(\mathcal{F}_2)$ and $(p(t''), q(t'')) \in \mathcal{F}_2(\mathcal{F}_1)$, there exists $\lambda \in (0, 1)$ such that $(p(\lambda t' + (1 - \lambda)t''), q(\lambda t' + (1 - \lambda)t'')) \in \mathcal{F}_3 \cup \mathcal{F}_4$. Similarly, $\mathcal{F}_4$ separates $\mathcal{F}_3$ from $\mathcal{F}_1$ and $\mathcal{F}_2$.

We now have sufficient information to draw the phase diagrams corresponding to the system (1.3)-(1.4) for all the three possible cases (i) $\alpha < \underline{\alpha}(L, R)$, (ii) $\alpha > \overline{\alpha}(L, R)$ and (iii) $\alpha \in (\underline{\alpha}(L, R), \overline{\alpha}(L, R))$. An example for each of these cases is given in Figures 1, 2 and 3, respectively, for a representative parameter configuration corresponding to each case. In all three plots the solid line represents $\overline{p}(\cdot)$ and the dashed line $\overline{q}(\cdot)$. The parameter values used for the plots are (i) $L = 3/5$, $R = 3/4$ and $\alpha = 1/2$, (ii) $L = 3/4$, $R = 3/5$ and $\alpha = 1/2$ and (iii) $L = R = 3/4$ and $\alpha = 1/2$. Even though these plots are for specific parameter configurations, the qualitative properties of the phase diagram are identical for each of the cases $\alpha < \underline{\alpha}(L, R)$, $\alpha > \overline{\alpha}(L, R)$ and $\alpha \in (\underline{\alpha}(L, R), \overline{\alpha}(L, R))$.

1.11 Appendix E: Proof of Proposition 1.5

Here we provide the lemmata used to prove Proposition 1.5. Define $[p^{-1}] : [0, 1] \to [0, 1]$ with $[p^{-1}](p) := \{q : \overline{p}(q) = p\}$ for $p \in [0, 1)$ and $[p^{-1}](1) = \max\{q \in [0, 1] : \overline{p}(q) = 1\}$. Define $[q^{-1}] : [0, 1] \to [0, 1]$ with $[q^{-1}](q) = \{p : \overline{q}(p) = q\}$ for $q \in [0, 1)$ and $[q^{-1}](1) = \max\{p \in [0, 1] : \overline{q}(p) = 1\}$. $[p^{-1}]$ and $[q^{-1}]$ are continuous, decreasing and concave.

Remark 11. *(i)* $\alpha_A > (<)\overline{\alpha}_A$ *if and only if* $\overline{p}'(0) < (>)[q^{-1}]'(0)$. *(ii)* $\alpha_B > (<)\overline{\alpha}_B$

if and only if $\overline{q}'(0) < (>) [p^{-1}]'(0)$.

Proof. (i) According to a well known result $[q^{-1}]'(0) = \frac{1}{\overline{q}'(q^{-1}(0))} = \frac{1}{\overline{q}'(1)}$. Next, $\overline{p}'(0) = \frac{\alpha_B R + R - 1}{R(1-\alpha_B)}$ and $\frac{1}{\overline{q}'(1)} = \frac{(1-L)(1-\alpha_A)}{\alpha_A(1-L)-L}$ and $\frac{\alpha_B R+R-1}{R(1-\alpha_B)} < \frac{(1-L)(1-\alpha_A)}{\alpha_A(1-L)-L}$ can be written $\alpha_A > \frac{L-R+\alpha_B R(1-2L)}{(1-L)(1-2R)} = \overline{\alpha}_A$. Analogous calculations hold for (ii). □

The following results follow from straightforward calculus:

Remark 12. $\overline{q}(\overline{p}(q)) - q = 0$ *is a polynomial equation of degree* 4, *and consequently has at most* 4 *different solutions in* $q \in (-\infty, \infty)$.

Remark 13. $\overline{q}(p)$ *has the following properties: (i)* $\overline{q}(0) = 1$ *and* $\overline{q}(1) = 0$, *(ii)* $\lim_{p\to\infty} \overline{q}(p) = -\lim_{p\to-\infty} \overline{q}(p) = \lim_{p\to\infty} \frac{\alpha_A(2L-1)-L}{(1-\alpha_A)(1-2L)} + \frac{\alpha_A}{1-\alpha_A} p$, *(iii)* $\overline{q}(p)$ *is discontinuous only at* $\frac{L}{2L-1}$. *In particular* $\lim_{p\to(\frac{L}{2L-1})^-} \overline{q}(p) = -\infty$ *and* $\lim_{p\to(\frac{L}{2L-1})^+} \overline{q}(p) = \infty$ *and (iv)* $\overline{q}(p)$ *has two local extrema, a local maximum at some* $p < 1$ *and a local minimum at some* $p > L/(2L-1)$.

The following Lemma shows that one of the four solutions to $\overline{q}(\overline{p}(q)) - q = 0$ occur outside of $[0, 1]$, which implies that if there is an internal rest point it is unique.

Lemma 1.10. *If* $\alpha_A \neq 1 - \alpha_B$, *then there is some* $\hat{q} \notin [0, 1]$ *such that* $\overline{q}(p(\hat{q})) - \hat{q} = 0$ *with* $\overline{p}(\hat{q}) \in (\frac{L}{2L-1}, \infty) \cup (0, -\infty)$.

Proof. (i) Consider the case $\frac{\alpha_A}{1-\alpha_A} < \frac{1-\alpha_B}{\alpha_B}$. Consider $\overline{q}(p)$ restricted to $(\frac{L}{2L-1}, \infty)$ and $\overline{p}(q)$ restricted to $(\frac{R}{2R-1}, \infty)$. Both $\overline{q}(p)$ and $\overline{p}(q)$ are continuous on $(\frac{1}{2L-1}, \infty)$ and $(\frac{R}{2R-1}, \infty)$ respectively. Since $\lim_{p\to\infty} \overline{q}(p) = \infty$ and $\lim_{q\to(\frac{R}{2R-1})^+} \overline{p}(q) = \infty$ there is a point $q' \in (\frac{R}{2R-1}, \infty)$ (close to $\frac{R}{2R-1}$) such that $q' < \overline{q}(\overline{p}(q'))$, which means that $(\overline{p}(q'), q')$ is in the subgraph of $\overline{q}(p)$. Next, given that $\frac{\alpha_A}{1-\alpha_A} < \frac{1-\alpha_B}{\alpha_B}$ there is some $q'' \in (\frac{R}{2R-1}, \infty)$ (sufficiently large) such that $\overline{p}(q'') \simeq \frac{\alpha_B(2R-1)-R}{(1-\alpha_B)(1-2R)} + \frac{\alpha_B}{1-\alpha_B} q''$, $\overline{q}(p(q'')) \simeq \frac{\alpha_A(2L-1)-L}{(1-\alpha_A)(1-2L)} + \frac{\alpha_A}{1-\alpha_A} \overline{p}(q'')$ and $q'' \simeq -\frac{(1-\alpha_B)(1-2R)}{\alpha_B(2R-1)-R} \frac{1-\alpha_B}{\alpha_B} + \frac{1-\alpha_B}{\alpha_B} \overline{p}(q'') >$

$\overline{q}(\overline{p}(q''))$. This means that $(\overline{p}(q''), q'')$ is in the epigraph of $\overline{q}(p)$. Since $\overline{p}(q)$ is continuous on $(\frac{R}{2R-1}, \infty)$, since $\lim_{p\to\infty} \overline{q}(p) = \lim_{p\to(\frac{L}{2L-1})^+} \overline{q}(p) = \infty$ and $\overline{q}(p)$ is continuous on $(\frac{L}{2L-1}, \infty)$, and since $(\overline{p}(q'), q')$ is in the subgraph of $\overline{q}(p)$ and $(\overline{p}(q''), q'')$ is in the epigraph of $\overline{q}(p)$, with $q', q'' \in (\frac{R}{2R-1}, \infty)$, there must be a point $(\overline{p}(\hat{q}), \hat{q})$ with $\hat{q} \in (\frac{R}{2R-1}, \infty)$ that is both in the subgraph and in the epigraph of $\overline{q}(p)$. Hence $\overline{q}(\overline{p}(\hat{q})) = \hat{q}$ for some $\hat{q} \notin [0,1]$, at which $p(\hat{q}) \in (\frac{L}{2L-1}, \infty)$.

Consider $\frac{\alpha_A}{1-\alpha_A} > \frac{1-\alpha_B}{\alpha_B}$. Consider $\overline{q}(p)$ and $\overline{p}(q)$ restricted to $[0, -\infty)$. Both $\overline{q}(p)$ and $\overline{p}(q)$ are continuous on $[0, -\infty)$. The point $(\overline{p}(q'), q')$, with $q' \in [0, -\infty)$ such that $\overline{p}(q') = 0$ is in the subgraph of $\overline{q}(p)$, since $\overline{q}(p(q')) = 1 > q'$.Next, given that $\frac{\alpha_A}{1-\alpha_A} > \frac{1-\alpha_B}{\alpha_B}$ there is some $q'' \in (0, -\infty)$ (sufficiently small) such that $\overline{p}(q'') \simeq \frac{\alpha_B(2R-1)-R}{(1-\alpha_B)(1-2R)} + \frac{\alpha_B}{1-\alpha_B} q''$, $\overline{q}(\overline{p}(q'')) \simeq \frac{\alpha_A(2L-1)-L}{(1-\alpha_A)(1-2L)} + \frac{\alpha_A}{1-\alpha_A}\overline{p}(q'')$ and $q'' \simeq -\frac{(1-\alpha_B)(1-2R)}{\alpha_B(2R-1)-R}\frac{1-\alpha_B}{\alpha_B} + \frac{1-\alpha_B}{\alpha_B}\overline{p}(q'') > \overline{q}(\overline{p}(q''))$. This means that $(\overline{p}(q''), q'')$ is in the epigraph of $\overline{q}(p)$.Since $\overline{p}(q)$ and $\overline{q}(p)$ are continuous on $[0, \infty)$, and since $(\overline{p}(q'), q')$ is in the subgraph of $\overline{q}(p)$ and $(\overline{p}(q''), q'')$ is in the epigraph of $\overline{q}(p)$, there must be a point $(\overline{p}(\hat{q}), \hat{q})$ with $\hat{q} \in (0, -\infty)$ that is both in the subgraph and in the epigraph of $\overline{q}(p)$. Hence $\overline{q}(\overline{p}(\hat{q})) = \hat{q}$ for some $\hat{q} \in (0, -\infty)$ at which $p(\hat{q}) \in (0, -\infty)$.

Hence, $q(p(\hat{q})) - \hat{q} = 0$ for some $\hat{q} \in (0, -\infty) \cup (\frac{R}{2R-1}, \infty)$. □

Lemma 1.11. *(i) If $\alpha_A > \overline{\alpha}_A$ or $\alpha_B > \overline{\alpha}_B$, then there is no internal rest point. (ii) If $\alpha_A < \overline{\alpha}_A$ and $\alpha_B < \overline{\alpha}_B$, then there is a unique internal rest point.*

Proof. (i) Suppose $\alpha_A > \overline{\alpha}_A$. Then by Remark 11 $\overline{p}'(0) < [q^{-1}]'(0)$. Define $[\widetilde{q}^{-1}] : [0, -\infty) \to \mathbb{R}$ with $[\widetilde{q}^{-1}](q) := \{p : \overline{q}(p) = q\}$ for $q \in [0, -\infty)$. Both $\widetilde{q}^{-1}(q)$ and $\overline{p}(q)$ are continuous on $[0, -\infty)$. Since $\overline{p}'(0) < [\widetilde{q}^{-1}]'(0)$ there is some $q' < 0$ (close to 0) such that $\overline{p}(q') > [\widetilde{q}^{-1}](q')$. Next, $\lim_{q\to\infty} \overline{p}(q) = -\infty$ while $\lim_{q\to\infty} [\widetilde{q}^{-1}](q) = \frac{L}{2L-1}$. This means that there is some q'' (sufficiently large) such that $\overline{p}(\underline{q}) > [\widetilde{q}^{-1}](\underline{q})$. Hence, there is some $\hat{q} \in (0, -\infty)$ such that $\overline{p}(\hat{q}) = [\widetilde{q}^{-1}](\hat{q})$ and $1 < p(\hat{q}) < \frac{L}{2L-1}$. Since there are at most four solutions to $\overline{q}(p(p)) - p = 0$ and there is one with $\overline{p}(q) \in (\frac{L}{2L-1}, \infty) \cup (0, -\infty)$ and one such that $\overline{p}(q) \in (1, \frac{L}{2L-1})$ there is no solution in $(0, 1)$. An analogous argument holds if $\alpha_B > \overline{\alpha}_B$.

(ii) Suppose. $\alpha_A < \overline{\alpha}_A$ or $\alpha_B < \overline{\alpha}_B$. By Remark 11 $\overline{q}'(0) > [p^{-1}]'(0)$, which means that there is some $p' \in (0,1)$ (close to 0) such that $\overline{q}(p') > [p]^{-1}(p')$. By Remark 11 $\overline{p}'(0) > [q^{-1}]'(0)$, which means that there is some $p'' \in (0,1)$ (close to 1) such that $\overline{q}(p'') < [p]^{-1}(p'')$. Since $[p]^{-1}(p)$ and $\overline{q}(p)$ are continuous on $(0,1)$ there is some $\hat{p} \in (0,1)$ such that $[p]^{-1}(\hat{p}) = \overline{q}(\hat{p})$ and hence there is an internal rest point. $\square$

Lemma 1.11 establishes, together with Lemma 1.1, that if we ignore the cases $\alpha_A \neq \overline{\alpha}_A$ and $\alpha_B \neq \overline{\alpha}_B$, then there is an internal rest point if and only if neither $(1,0)$ nor $(0,1)$ are asymptotically stable.

Remark 14. *(i)* $\hat{p}_1(\alpha_A, \alpha_B) > 0$, *(ii)* $\hat{p}_2(\alpha_A, \alpha_B) < 0$, *(iii)* $\hat{q}_1(\alpha_B, \alpha_A) > 0$ *and (iv)* $\hat{q}_2(\alpha_B, \alpha_A) < 0$.

Proof. $\hat{q}(\alpha_B, \alpha_A)$ is defined by $\overline{q}(p(\hat{q})) - \hat{q} = 0$. To establish (iii), we differentiate implicitly and obtain $\frac{\partial \hat{q}}{\partial \alpha_B} = -\frac{q'(\overline{p}(\hat{q}))\frac{\partial p(\hat{q})}{\partial \alpha_B}}{q'(\overline{p}(\hat{q})\overline{p}'(\hat{q})-1}$. Consider first the denominator. It holds that $[p^{-1}]'(\hat{p}) > \overline{q}'(\overline{p}(\hat{q}))$. This means $\frac{1}{\overline{p}'(\hat{q})} > \overline{q}'(\overline{p}(\hat{q}))$, or $1 < \overline{q}'(\overline{p}(\hat{q}))\overline{p}'(\hat{q})$. Therefore, the denominator is positive. Now consider the denominator. Note that $\frac{\partial \overline{p}(\hat{q})}{\partial \alpha_B} > 0$ and $\overline{q}'(\overline{p}(\hat{q})) < 0$. Hence, the numerator is negative, so $\frac{\partial \hat{q}}{\partial \alpha_B} > 0$. Proceeding in the same way with (iv) we obtain $\frac{\partial \hat{q}}{\partial \alpha_B} = -\frac{\frac{\partial q(\overline{p}(\hat{q}))}{\partial \alpha_A}}{\overline{q}'(\overline{p}(\hat{q})\overline{p}'(\hat{q})-1} < 0$. Analogous arguments apply to (i) and (ii). $\square$

Remark 15. *If* $\alpha_B > \frac{1-R}{R}$, *then* $\lim_{\alpha_A \to 1} \hat{p}(\alpha_A, \alpha_B) = 1$ *and if* $\alpha_A > \frac{1-L}{L}$, *then* $\lim_{\alpha_B \to 1} \hat{q}(\alpha_B, \alpha_A) = 1$.

Proof. Consider a value of α_A large enough such that there is some $p' \in (0,1)$ with $\dot{p}(p', 1) = 0$. Then since $\overline{q}(p)$ is concave and passes through $(p', 1)$ and $(1, 0)$ it holds that $\hat{p}(\alpha_A, \alpha_B) > p'$. Next, $\dot{p}(p', 1) = \alpha_A p'(1-p')(2L-1) - (1-\alpha_A)p'(L-1) = 0$ means $p' = \frac{\alpha_A L + L - 1}{\alpha_A(2L-1)}$, which approaches 1 as $\alpha_A \to 1$. Since $\hat{p}(\alpha_A, \alpha_B) > p'$ it must then hold that $\lim_{\alpha_A \to 1} \hat{p}(\alpha_A, \alpha_B) = 1$. An analogous argument holds for $\hat{q}(\alpha_B, \alpha_A)$. $\square$

Bibliography

[1] Alos-Ferrer, C. and Schlag, K.: Imitation and learning. Ch. 11, The Handbook of Rational and Social Choice, Anand, P.; Pattanaik, P.; Puppe, C., (Eds.) pp. 271-298(28), Oxford University Press (2009).

[2] Bayer, P. Ross, S., Topa, G.: "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes," J. Polit. Economy, 116:6, pp. 1150-1196 (2008)

[3] Conley, T., Udry, C.: Learning About a New Technology: Pineapple in Ghana. American Economic Review 100, 35–69 (2010)

[4] Cubitt, R., Sugden, R.: The selection of preferences through imitation. Rev. Econ. Stud. 65, 761-771 (1998)

[5] Currarini, S., Jackson, M., Pin, P.: An Economic Model of Friendship: Homophily, Minorities, and Segregation. Econometrica 77, 1003 - 1045, (2009)

[6] Dixon, R.: Hybrid Corn Revisited. Econometrica 48, 1451–1461 (1980)

[7] Ellison, G., Fudenberg, D.: Rules of thumb for social learning. J. Polit. Economy 101, 612-643 (1993)

[8] Ellison, G., Fudenberg, D.: Word of mouth communication and social learning. Quart. J. Econ. 110, 93-125 (1995)

[9] Festinger, L.: A theory of social comparison processes. Human Relations, 7, 117-140 (1954)

[10] Foster, A., Rosenzweig, M.: Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. J. Polit. Economy, 103, 1176–1209 (1995)

[11] Geroski, P.: Models of Technology Diffusion." Research Policy 29, 603–25 (2000)

[12] Griliches, Z.: Hybrid Corn: An Exploration of the Economics of Technological Change. Econometrica, 25: 501–22. (1957)

[13] Henrich, J.: Cultural Transmission and the Diffusion of Innovations: Adoption Dynamics Indicate that Biased Cultural Transmission is the Predominant Force in Behavioral Change." American Anthropologist 103, 992–1013 (2001)

[14] Hofbauer, J., Sigmund, K.: Evolutionary Games and Population Dynamics, Cambridge University Press, 1998.

[15] Manski, C.F.: Economic analysis of social interactions. J. Econ Perspectives, 14, pp. 115-136 (2000).

[16] Munshi, K.: Social learning in a heterogeneous population. J. Dev. Econ. 73, 185-213 (2004)

[17] Oyarzun, C., Ruf, J.: Monotone imitation. Econ. Theory 41, 411-441 (2009).

[18] Persico, N.: Information acquisition in auctions. Econometrica, 68, pp. 135-148 (2000)

[19] Rogers, E.: Diffusion of Innovations. 5th ed. New York: Free Press (2003)

[20] Ryan, B, Gross, N.: The Diffusion of Hybrid Corn in Two Iowa Communities. Rural Sociology 8, 15–24 (1943)

[21] Sandholm, W. Population Games and Evolutionary Dynamics. Cambridge MA: MIT press (2010)

[22] Santos-Pinto, L., Sobel, J.: A Model of Positive Self-Image in Subjective Assessments, American Economic Review 95, 1386-1402 (2005)

[23] Schlag, K.: Why imitate, and if so, how? A bounded rational approach to multi-armed bandits. J. Econ. Theory 78, 130-156 (1998)

[24] Schlag, K.: Which one should I imitate. J. Math. Econ. 31, 493-522 (1999)

[25] Sorensen, A.: Social learning and health plan choice. RAND J. of Econ. 37, pp. 929-945 (2006)

[26] Suls, J., Martin, R., Wheeler, L.: Three Kinds of Opinion Comparison: The Triadic Model. Personality and Social Psychology Review 4, 219-237 (2000)

[27] Suri, T.: Selection and Comparative Advantage in Technology Adoption. Econometrica 79, 159-209 (2011)

[28] Vega-Redondo, F.: The evolution of Walrasian behavior. Econometrica 65, 375-384 (1997)

[29] Young, P.: Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. American Economic Review 99, 1899–1924 (2009)

Chapter 2

Imitation in Cournot Oligopolies with Multiple Markets

2.1 Introduction

A fundamental assumption in many economic models is that agents arrive at their decisions by maximizing an objective function. In Cournot oligopolies, this translates into the typical assumption of profit maximization. However, in recent years a literature has emerged that studies quantity decisions based on imitation of successful behavior, rather than on maximization of profits. There are several reasons to do this. Oligopolies are complex situations, which means that it is likely that firms sometimes make decisions through processes that are cognitively less demanding than profit maximization (e.g. Conlisk 1996 or Gigerenzer, Todd and the ABC Research Group 1999). At the same time, imitation of successful behavior is an intuitive and cognitively simple heuristic, and is often observed in everyday life. Sometimes, it eventually leads an entire population to the optimal choice (Schlag 1998). Several experimental papers have also documented imitative behavior in oligopoly games (Huck, Normann and Oechssler 1999, 2000; Offerman, Potters and Sonnemans 2002; Apesteguia and Selten 2005; Apesteguia, Huck and Oechssler 2007, 2008) as well as in other situations[1].

[1] For example, Pingle and Day (1996) focus on imitation as a means to reduce decision costs and Offerman and Schotter (2009) study imitation in situations characterized by exogenous uncertainty.

However, in oligopolies there is a tendency for imitative behavior to lead to very competitive outcomes and corresponding low profits, even if the firms imitate only firms which obtained *high profits* (Vega-Redondo 1997). The reason is that firms competing aggressively tend to do better and therefore also to be imitated, which leads to competitive behavior in the market. On the other hand, this tendency is reduced if firms sometimes imitate non-competitors, acting in other markets. In this case, firms in markets displaying more cooperative behavior obtain higher profits. This leads more cooperative strategies to sometimes be imitated (Apesteguia, Huck and Oechssler 2007). This suggests that the extent to which the firms have information about firms in other markets is important for the outcome when quantities are chosen by imitation. Yet, there is not much literature exploring how different assumptions with respect to the availability of such information affect the outcome. This paper contributes to the literature on imitation in Cournot games by thoroughly studying the countervailing tendencies that arise when firms observe and imitate both competitors and non-competitors. The key parameter under study is the amount of information available about firms in other markets. The general conclusion is that more information about firms in other markets leads to less competitive outcomes.

I assume that there is a set of identical and separated markets. In each period each firm observes the quantities and profits obtained in the previous period by the firms in the own market, as well as those of a sample of firms from the other markets. They then choose a quantity by using a decision rule. Either, they choose the quantity that generated the highest profit in the sample. This decision rule is referred to as *Imitate the Best Max* (henceforth, IBM). IBM has been studied in a number of different contexts and is often motivated by its simplicity and by the natural salience of high payoffs.[2] However, when there are several markets, a quantity can perform well in one market and poorly in another. In this case it is not clear that firms would find it appealing to follow IBM and imitate such a quantity. The

[2] See Alós-Ferrer 2004 for a discussion of IBM. It has been studied, e.g., in Ellison and Fudenberg 1993; Vega-Redondo 1997; Tanaka 1999, 2000; Alós-Ferrer and Ania 2005.

second decision rule that I consider takes this into account by computing the average profit of each observed quantity and choosing the quantity that generated the highest average profit. This decision rule is referred to as *Imitate the Best Average* (henceforth IBA).[3] Firms also sometimes experiment and choose a quantity at random. I analyze the resulting stochastic process by using the techniques developed by Young (1993) to characterize the *stochastically stable states*, i.e. the support of the limiting distribution of the stochastic process. The stochastically stable states constitute the long run prediction of the model.

Three different informational settings are considered. In a benchmark case, firms observe all firms in all markets. If firms use IBM, all quantities between the Cournot and *Walrasian* quantities appear in the long run[4] (Proposition 2.1). Intuitively, quantities closer to the Walrasian quantity tend to perform the best in each market, whereas markets closer to the Cournot quantity tend to perform better than other markets. This creates countervailing tendencies in the dynamics which balance out in the long run. If firms instead use IBA, then there is a unique stochastically stable state, in which all firms produce the same quantity (Proposition 2.5). This quantity, which I denote q^s, lies strictly between the Cournot and Walrasian quantities. Intuitively, if all firms produce q^s, an experimentation toward the Walrasian quantity increases the profits with respect to the firms in the own market. However, this is outweighed by the decrease in profits in absolute terms (i.e. in comparison to firms in other markets) in the computation of average profits. The converse holds for experimentation toward the Cournot quantity. This means that q^s is in some sense stable. It turns out that it is the only quantity with these properties. Moreover, q^s decreases in the number of markets and increases in the number of firms per market. This is a prediction that is suitable for experimental testing. It is also shown that q^s corresponds to the unique symmetric Nash equilibrium of a game in which firms are concerned about profits in both absolute and relative terms (Proposition 8).

[3]IBA has been studied, e.g., by Eshel, Samuelson and Shaked 1998; Jun and Sethi 2007; Bergin and Bernhardt 2009; Apesteguia et al. 2007; and Mengel 2009.

[4]I follow Vega-Redondo 1997 and refer to the quantity that each firm produces in a symmetric perfectly competitive equilibrium as the *Walrasian* quantity.

In the second informational setting, firms observe markets where the aggregate quantity produced is sufficiently close to the aggregate quantity in the own market. The idea is that imitation intuitively makes more sense if the sampled firms are in a similar situation as oneself. In this way, the extent to which firms imitate across markets can be continuously varied, from the benchmark case studied here, to the single-market case analyzed by Vega-Redondo (1997), including all intermediate cases. Moreover, the willingness to imitate across markets is endogenous, since it is sensitive to ex-post differences between the markets. In Proposition 2.9, it is shown that if firms are sufficiently insensitive to differences in the aggregate, the model behaves as in the benchmark case. If firms instead are very sensitive, markets evolve independently and we obtain the perfectly competitive outcome (as in Vega-Redondo (1997)). For intermediate cases firms produce quantities in an interval with lower bound strictly between the Cournot and Walrasian quantity and upper bound at the Walrasian quantity. The lower bound is increasing in firms' sensitivity to differences in the aggregate. Hence, less competitive outcomes are obtained when firms are less willing to imitate across markets.

In the third informational setting the markets are arranged around a circle and firms observe some of the neighboring markets. For example, this can reflect a situation in which firms' geographical locations prevent them from observing all the remaining markets. As long as there is some positive inertia (in the sense that not all firms necessarily adjust quantities in each period), the results from the benchmark case are robust to this setting (Proposition 2.11 and 12). This means that if firms use IBA, the outcome becomes less competitive as firms observe a larger set of markets. A conclusion that holds across the different informational settings is that more information tends to lead to less competitive outcomes. This contrasts with the conclusion of Huck et. al. (1999, 2000), that more information *about the firms in the own market* leads to more competitive results.

This paper is related to the seminal paper of Vega-Redondo (1997), who showed that if decisions are made through imitation, a single market converges to the perfectly competitive outcome. There is also a close relationship with Apesteguia, et.

al. (2007), who study a model in which there are several markets and firms sometimes imitate across markets. Apesteguia et. al. (2007) were the first to show that in this case imitation need not lead to competitive outcomes. This is shown in a setting with linear demand, zero costs, a capacity constraint at the Walrasian output and markets that are remixed from one period to another. The present paper reinforces the conclusion of Apesteguia et. al. (2007) and complements their analysis in several ways. First, by considering information structures different from those in Apesteguia et. al. (2007) it is shown that outcomes gradually becomes less competitive as firms' tendency to imitate across markets gradually increases. This analysis is not present in Apesteguia et. al. (2007). Second, I consider more general demand and cost functions, and relax the assumption of a capacity constraint at the Walrasian quantity and that markets are remixed. The general conclusion that imitation across markets reduces competition, first uncovered by Apesteguia et. al. (2007), is robust to this more general setting. However, the specific results obtained here are very different from those of Apesteguia et. al. (2007). For example, here the benchmark setting predicts all quantities between the Cournot and Walrasian quantity for IBM and a unique quantity in this interval for IBA. In contrast, in a similar informational setting Apesteguia et. al. (2007) predict the Walrasian and Cournot outcome for IBM and IBA, respectively. Hence, the specific results obtained by Apesteguia et. al. (2007) indeed depend on the specific set of assumptions used by these authors.

This paper is also related to Alós-Ferrer (2004) and Bergin and Berghardt (2009). In these papers, there is a single market but firms remember past quantities and profits. The relationship arises since the memory model can be seen as a multimarket model, in which the additional markets exist in the memories of the firms. I discuss this relationship in Appendix C and use it in section 3.1 to prove Proposition 2.1. Finally, this paper is related to the various extensions on Vega-Redondo (1997), for example to asymmetric oligopolies (Tanaka, 1999), industries of differentiated goods (Tanaka, 2000), more general technical conditions (Schenk-Hoppé, 2000) more general classes of games (Alós-Ferrer and Ania, 2005) and industries with both optimizers and imitators (Schipper, 2008).

The model is presented in Section 2. Section 3 contains the results for IBM for the benchmark informational setting and Section 4 contains those for IBA. Section 5 analyzes the second and third informational settings. Section 6 concludes.

2.2 The Model

2.2.1 Market Structure

Consider a population of kn firms consisting of $k \geq 2$ disjoint groups with $n \geq 2$ firms in each. Let $K = \{1, 2, ..., k\}$ denote the set of groups and $N = \{1, 2, ..., n\}$ denote the firms in a generic group. A group is thus identified by an index $j \in K$ and an individual firm is identified by a double index $ij \in N \times K$, where j indicates the group and i is the identity of the individual firm within the group. The model proceeds in discrete time. At each point in time $t = \{0, 1, ...\}$ each firm $ij \in N \times K$ produces a quantity $q_{ij}(t) \in \Gamma$ at cost $C(q_{ij}(t))$[5], where $\Gamma = \{0, \delta, 2\delta, ..., v\delta\}$ is a common finite grid and the step size δ can be arbitrarily small[6]. In Apesteguia et. al. (2007) $v\delta = q^w$ (q^w is defined below), which turns out to be important for the results they derive. Here we assume that $v\delta$ is "very large", in the sense that there will not be a capacity constraint that affect the dynamics of the system. The good is sold in group-specific, completely isolated markets, facing a demand represented by a function $P(Q_j)$, where $Q_j = \sum_{i \in N} q_{ij}$, $\forall j \in K$. The profit of firm ij can then be written $\pi(q_{ij}, Q_j) = P(Q_j)q_{ij} - C(q_{ij})$. A firm thus competes only with firms belonging to the same group, introducing a sharp division between local, neighboring firms and those operating in other markets. For tractability we work with a well behaved Cournot oligopoly, i.e. we assume:

Assumption 1: $C(q)$ is twice continuously differentiable, $C'(q) > 0$ and

[5] In order to keep notation light, in what follows we suppress time indexes as long as this does not create any ambiguity.

[6] This is a technical assumptions which ensures that we will be working with a finite markov chain, which allows us to use the techniques of Young (1993) when analyzing the dynamics of this system.

$C'''(q) \geq 0$.

Assumption 2: $P(Q)$ is twice continuously differentiable, $P'(Q) < 0$ and $P''(Q) \leq 0$.

Assumption 3: $P(0) > 0$ and there is a quantity $\underline{Q}$ such that $P(Q) = 0$, $\forall Q \geq \underline{Q}$.

Next we proceed to define the following quantities:

Definition 2.1. *The symmetric Walrasian output* q^w *is a quantity such that* $P(nq^w)q^w - C(q^w) \geq P(nq^w)q' - C(q')$, $\forall q' \in \Gamma$.

Definition 2.2. *The symmetric Cournot output* q^c *is a quantity such that* $P(nq^c)q^c - C(q^c) \geq P((n-1)q^c + q')q' - C(q')$, $\forall q' \in \Gamma$.

As usual the symmetric Walrasian output is a quantity that maximizes profits taking the price as given, whereas the symmetric Cournot output maximizes profit taking the quantities produced by the remaining firms as given. Assumptions 1-3 guarantee that unique symmetric Walrasian and Cournot equilibria exist. We assume for simplicity that q^w and q^c belong to the grid Γ.

2.2.2 Decision and Dynamics

Firms do not necessarily decide upon quantities in all periods. We assume that in each period firms are picked independently with identical probability $\mu \in (0, 1]$ to revise the strategy. When $\mu \in (0, 1)$ we say that decisions are revised with inertia. Inertia reflects both an unwillingness to adjust the quantity too often and the possibility that firms do not adjust their quantities in a perfectly coordinated way. The main results of this paper do not depend on the presence or absence of inertia and the proofs we provide hold for both cases.

We will consider two different decision rules. The first one is called *Imitate the Best Max* (IBM). Under this rule an agent adjusting his strategy[7] in t observes the quantities and profits of all firms in $t-1$. He then assigns positive probabilities according to some probability distribution of imitating all those quantities that generated the highest payoff in the population in the previous period. Formally, the agent assigns positive probabilities to all quantities in

$$B(t)^{IBM} = \underset{q_{ij}(t-1):ij \in N\times K}{\arg\max} \{\pi(q_{ij}(t-1), Q_j(t-1)). \quad (2.1)$$

IBM simply prescribes imitating the most successful firm observed. This rule has received a great deal of attention in the literature and is considered in for example Vega-Redondo (1997), Tanaka (1999, 2000), Alós-Ferrer (2004) and Alós-Ferrer and Ania (2005). It has found experimental support in Apesteguia et. al. (2007).

However, as argued by Apesteguia et al. (2007), in a multi-market framework the same strategy is likely to obtain different payoffs in different markets. It is then not clear that a firm would find it attractive to imitate the quantity that generated the highest maximum payoff, since this same quantity may have performed poorly in other markets. A more prudent imitator may want to take into account this information. One way of doing this is by computing the average payoffs of the observed quantities and copy the quantity that rendered the highest average payoff. We refer to this rule as *Imitate the Best Average* (IBA). Let $\{q_{ij}(t-1)\} = \{i'j' \in N \times K : q_{i'j'}(t-1) = q_{ij}(t-1)\}$. Then, an agent using IBA assigns positive probabilities to all quantities in

$$B(t)^{IBA} = \underset{q_{ij}(t-1):ij \in N\times K}{\arg\max} \frac{1}{|\{q_{ij}(t-1)\}|} \sum_{\{q_{ij}(t-1)\}} \pi(q_{i'j'}(t-1), Q_{j'}(t-1)). \quad (2.2)$$

IBA has also received attention in the literature, but in oligopoly models far less than IBM. A reason for this is that when there is only one market, as in most existing papers, the rules are equivalent and there is thus no point in talking about

[7] By "strategy" we mean the choice of a quantity.

averages. IBA becomes suitable when there is localness in some sense in the model, such as in the present paper. It has been considered in for example Eshel, Samuelson and Shaked (1998), Jun and Sethi (2005), Bergin and Bernhardt (2009), Apesteguia et al. (2007) and Mengel (2009).

The model outlined so far defines a finite Markov chain. A state of this process is a vector of the quantities produced by the different firm, denoted by $(q_{11}, q_{21}, ..., q_{n1};$ $q_{12}, q_{22}, ..., q_{n2}; ...; q_{1k}, q_{2k}, ..., q_{nk})$. The state space is Γ^{kn}. The transition probabilities are determined by the inertia parameter μ and the imitation rule. We follow the convention in the literature and refer to the process described so far as the *unperturbed process*. We say that a state x such that $P_{x,y}^{(m)} = 0$, $\forall y \in \Gamma^{kn} \backslash \{x\}$, $\forall m \in \mathbb{N}$ is an absorbing state, where $P_{x,y}^{(m)}$ is the probability of reaching state y in m steps, starting from x. An absorbing state is then a state that once entered is left with probability zero. In the unperturbed process, any state in which all firms produce the same quantity is an absorbing state. Moreover, the fact that the imitation rule assigns positive probabilities to all quantities that generated the highest (or the highest average) payoff implies that only the states in which all firms produce the same quantity are absorbing. We denote by $\omega(q)$ the state in which all firms produce q and state what we just mentioned as a first lemma. Let $\Omega_M := \{\omega(q) : q \in \Gamma\}$[8], i.e. Ω_M is the set of states in which all firms produce the same quantity, and denote the set of absorbing states by Ω. Then, as in Vega-Redondo (1997):

Lemma 2.1. $\Omega = \Omega_M$.

Proof. As argued above. □

Next, following the common methodology in the literature, we incorporate an "error term" into the model, which captures deviations from the behavior prescribed by the imitation rule. Such deviations may be due to experimentation or an error on part of the firm. We assume that a firm follows the prescription of the imitation

[8]The "M" stands for "monomorphic", reflecting the fact that all firms use the same strategy.

dynamics with probability $1-\epsilon$ and that with probability ϵ it picks a quantity randomly from Γ according to some probability distribution over Γ with full support. These experimentations are identically and independently distributed across firms. With the error term a different Markov process results. We refer to this process as the *perturbed process.* Since an agent deviating from the decision rule picks any quantity with positive probability, a transition between any two states occurs with positive probability and the perturbed process is therefore ergodic. Then, by a standard result a unique limiting distribution exists which describes average behavior in the long run. We use the techniques developed by Freidlin and Wentzell (1988) and introduced into economics by Young (1993) and Kandori, Mailath and Rob (1993), to find the support of this limiting distribution as ϵ goes to zero, referred to as the *stochastically stable states,* which we will denote by Ω^{ss}. As noted by these authors, $\Omega^{ss} \subseteq \Omega$. This implies that we can focus our study on the absorbing states, which by lemma 1 coincides with the states in which all firms produce the same quantity. There are two different ways of finding the stochastically stable states. The first one involves the construction of a tree originating in a certain absorbing state, where the remaining absorbing states constitute the nodes, and the weights or costs of the edges is the minimum number of experimentations needed for a transition from one absorbing state to another. Finding the stochastically stable states involves finding the state with the minimum cost tree. (For this approach see Young (1993) and Kandori et al. (1993)). The other way is a shortcut due to Ellison (2000), who observes that a sufficient condition for stochastic stability of a state is that the minimum number of experimentations necessary to exit this state be smaller than the maximum number of experimentations needed to enter it. We sketch these results in Appendix A[9]. The interpretation of the stochastically stable states is that in the long run, the process will spend almost all of its time in these states.

[9] We refer the reader to Fudenberg and Levine (1998) for a treatment of these methods.

2.3 Imitate the best max

2.3.1 Stochastically Stable States

To characterize the stochastically stable states, it is of importance to determine the conditions under which experimentation leads the system with positive probability from one absorbing state to another. In the multi-market model, regardless of whether $\mu = 1$ or $\mu \in (0, 1)$, a single experimentation q' leads the system from $\omega(q)$ to $\omega(q')$ with positive probability if and only if

$$D(q, q') := \pi(q', (n-1)q + q') - \max\{\pi(q, (n-1)q + q'), \pi(q, nq)\} \geq 0. \qquad (2.3)$$

That is, an invading strategy must obtain the highest profit in the own market and a higher profit than the firms in other markets. In other words, the experimenter must achieve both a relative payoff improvement (outdoing the competition) and an absolute payoff improvement (increasing his own profit). If $n, k > 2$, two deviations facilitate such a transition if and only if $\pi(q', (n-2)q + q' + q'') - \max\{\pi(q, (n-2)q + q' + q''), \pi(q, nq)\} \geq 0$, where q'' is the second deviating quantity. In the special case of $k = 2$, if a single deviator is placed in each market it is sufficient that $\pi(q', (n-1)q + q') \geq \pi(q, (n-1)q + q')$, ($q'$ is better in relative terms). If $n = 2$, then $\pi(q', nq') \geq \pi(q, nq)$ is sufficient, i.e. any quantity that is "collusively" better makes a transition possible.

In order to find the stochastically stable states we will rely heavily on Alós-Ferrer (2004), who analyzes firms who use IBM in a single Cournot market in which firms recall past quantities and profits. The reason that we can use his results is that a single market model with memory closely resembles a multi-market model. Intuitively, in a single market model with memory, the additional, outside markets in a sense exist in the memory of agents.

Proposition 2.1. *When firms use IBM and there are k markets with n firms in each, the set of stochastically stable states is $\{\omega(q) : q \in [q^c, q^w]\}$[10].*

[10] The correct notation is $\{\omega(q) : q \in [q^c, q^w] \cap \Gamma\}$, but here and in what follows we suppress the

Proof. We will use the proofs of theorems 1 and 2 in Alós-Ferrer (2004). He shows :

(i) $D(q, q') < 0$ for all $q \in [q^c, q^w]$, $q' \neq q$.

(ii) For any $q \in \Gamma \backslash [q^c, q^w]$ there is some $q' \in [q^c, q^w]$ such that $D(q, q') > 0$.

(iii) For any $q \in [q^c, q^w)$ we have $\pi(q', (n-2)q + q') - \max\{\pi(q, (n-2)q + q'), \pi(q, nq)\} > 0$ for $q' \in (q, q^w]$.

(iv) For any $q \in [q^c, q^w]$ there is a $\varphi(q) > q^w$ such that $\pi(\varphi(q), (n-2)q + \varphi(q)) - \max\{\pi(q, (n-2)q + \varphi(q)), \pi(q, nq)\} > 0$ and at the same time there is some $q' < q$ such that $D(\varphi(q), q') > 0$ for all $[q', \varphi(q))$.

A combination of radius-coradius and tree-surgery arguments (see Appendix A) can now be used to show that $\{\omega(q) : q \in [q^c, q^w]\}$ is the set of stochastically stable states.

First, (i) means that $R(\{\omega(q) : q \in [q^c, q^w]\}) > 1$ and (ii) implies that $Cr(\{\omega(q) : q \in [q^c, q^w]\}) = 1$. Hence, $R(\{\omega(q) : q \in [q^c, q^w]\}) > Cr(\{\omega(q) : q \in [q^c, q^w]\}) = 1$, which means that the set of stochastically stable states is contained in $\{\omega(q) : q \in [q^c, q^w]\}$.

Next, (iii) implies that a $\omega(q^w)$-tree can be constructed in which all arrows exiting states in $\{\omega(q) : q \in [q^c, q^w)\}$ have cost 2, which means that $\omega(q^w)$ has minimum stochastic potential and is therefore stochastically stable.

Finally (iv) implies that for any state in $\{\omega(q) : q \in [q^c, q^w)\}$ a $\omega(q)$-tree can be constructed in which all states in $\{\omega(q) : q \in (q, q^w]\}$ have exiting arrows at cost 2. This is accomplished through sequences of transitions $\omega(q^i) \xrightarrow{2} \omega(\varphi(q^i)) \xrightarrow{1} \omega(q^i) \xrightarrow{2} \omega(\varphi(q^i)) \xrightarrow{1} \omega(q^{i+1}) \rightarrow ... \rightarrow \omega(q^{i+1})$, in which $q^i \in (q, q^w]$ and $q^i > q^{i+1} > q^{i+1}$... and so on ($\omega(q) \xrightarrow{x} \omega(q')$ denotes a transition from $\omega(q)$ to $\omega(q')$ through x experimentations). By (iii) the arrows exiting states in $\{\omega(q) : q \in [q^c, q)\}$ also have cost 2. The implication is that $\omega(q)$ has minimum stochastic potential and is therefore stochastically stable. Since this can be done for any $\{\omega(q) : q \in [q^c, q^w)\}$, all these states are stochastically stable. □

The set of stochastically stable states here corresponds precisely to those in Alós-

$\cap\Gamma$ part to lighten notation.

Ferrer (2004), which points to the close relationship between the single market model with memory and the multimarket model analyzed here. An important difference is that Alós-Ferrer's (2004) results hold only under an assumption of no inertia, whereas inertia is inconsequential in the setting analyzed here. In Appendix C we discuss a bit more closely the relationship between the model of Alós-Ferrer (2004) and the model analyzed here.

The intuition behind Proposition 2.1 is that quantities outside the interval $[q^c, q^w]$ are easily destabilized since experimentations that are better in both relative and absolute terms are available. Quantities within this interval are harder to destabilize, since an experimentation cannot be better both in relative and absolute terms. In fact, for any $q \in [q^c, q^w]$ a deviation $q' > (<) q$ is better (worse) in absolute terms, but worse (better) in relative terms. The result is that the entire interval $[q^c, q^w]$ is stochastically stable.

Proposition 2.1 should be contrasted with the corresponding result in Apesteguia et. al. (2007), who obtain that $\omega(q^w)$ is the unique stochastically stable state when firms use IBM. The difference arises mainly due to the capacity constraint they impose at q^w. As seen in the proof of Proposition 2.1, in order to leave $\omega(q^w)$ with two experimentations, one firm must experiment to a quantity $q' > q^w$. With the capacity constraint this is not possible. In fact, it can be shown that if we impose some capacity constraint $\overline{v\delta} \in [q^c, q^w]$ the unique stochastically stable state is $\omega(\overline{v\delta})$.

Proposition 2.1 can be related to the concept of finite population evolutionary stability (Schaffer (1988)). An evolutionary stable strategy according to the definition of Schaffer (1988), is a strategy q such that any single deviation from $\omega(q)$ leaves the deviator worse off than the incumbents. That is, if we consider a single population, q is evolutionary stable if and only if $\pi(q', (n-1)q' + q) \leq \pi(q, (n-1)q' + q)$ for any $q' \neq q$. If we extend this definition to the multi-market setting by requiring any deviator to be worse off than either the incumbents in his group or those in other groups, the above result implies that $[q^c, q^w]$ are evolutionary stable strategies, whereas $\Gamma \backslash [q^c, q^w]$ are not.

2.3.2 Maximization over Aspiration Levels

It is possible to give an alternative interpretation of the result obtained above. Schaffer (1988) shows that in the case of a single market, the evolutionary stable strategy can be obtained as a symmetric Nash equilibrium in a setting where firms maximize relative payoffs. A result of a similar flavor can be obtained here. Consider a single Cournot market as specified above. Let each firm $i \in N$ in each period t choose a strategy in

$$\begin{aligned} q_i^* &= \arg\max_{q_i}\{\pi(q_i, Q_{-i}(t-1)+q_i) - \\ &\quad \max\{K(\pi(q_{-i}, Q_{-i}(t-1)+q_i)), \pi(q_i(t-1), Q(t-1))\}. \end{aligned} \tag{2.4}$$

where $Q_{-i}(t-1) = \sum_{j \neq i} q_j(t-1)$, and $K(\pi(q_{-i}, Q_{-i}(t-1)+q_i))$ is a convex combination of the profits of the firms in $N \backslash \{i\}$. This means that the firms choose a strategy to myopically maximize profits over the maximum of two aspiration levels: the profits of the competition and the own profit obtained in the previous period. Note that if $K(\pi(q_{-i}, Q_{-i}(t-1)+q_i))$ were removed from (4), the decision rule is equivalent to one of *myopic best response* (see for example Vega-Redondo (2003)). We refer to the dynamic system implied by decision rule (4) as *Maximization over Aspiration levels*. We say that the system is in a symmetric rest point if $q_i^* = q_i(t-1) = q^*$ for all $i \in N$. Then:

Proposition 2.2. *The set of symmetric rest points of the system Maximization over Aspiration Levels is* $\{\omega(q) : q \in [q^c, q^w]\}$.

Proof. The quantity q^* is a symmetric rest point if and only if $q^* = \arg\max_{q_i}\{\pi(q_i, (n-1)q^* + q_i) - \max\{\pi(q^*, (n-1)q^* + q_i), \pi(q^*, nq^*)\} = \arg\max_{q'} D(q^*, q')$. We note that $D(q, q) = 0$ for all $q \in \Gamma$. Together with (i) in the proof of Proposition 2.1, this implies that $\arg\max_{q'} D(q, q') = q$ for all $q \in [q^c, q^w]$ and hence these quantities are symmetric rest points. Next, $D(q, q) = 0$ for all $q \in \Gamma$ together with (ii) implies that

$\arg\max_{q'} D(q, q') \neq q$ for $q \in \Gamma \backslash [q^c, q^w]$ and hence these quantities are not symmetric rest points. $\square$

In an analogous way it can also be obtained that $[q^c, q^w]$ is the set of symmetric Nash equilibria when firms maximize profits over the maximum of the profit obtained by firms in the same market and in other markets.

The prediction obtained with IBM is thus closely related to that obtained in a system in which firms are concerned about both outdoing the competition and increasing profits in absolute terms. An implication of the above decision rule is that a firm is unwilling to choose a strategy that improves absolute payoffs if this worsens its position with respect to the competition too much. On the other hand, a firm is not willing to produce a quantity that improves its position with respect to the remaining firms if by doing this its absolute payoffs decreases to a sufficient degree. The result is that quantities in $[q^c, q^w]$ become symmetric rest points. It turns out that at these points it is not possible to choose a quantity that improves both absolute payoffs and position with respect to the remaining firms.

As mentioned, IBM may be overly naive in a multimarket context, since a given strategy may render different profits in different markets and this is not taken into account by this rule. In the following section we study a rule that does take this into account.

2.4 Imitate the Best Average

Again, to analyze the dynamic properties of the system we focus on the absorbing states of the process - the states in which all firms produce the same quantity. We are interested in pinpointing the conditions under which experimentation makes transitions from one absorbing state to another occur with positive probability. Fortunately, the dynamics under IBA are "well behaved" and it is actually sufficient to consider single experimentations to find the stochastically stable state. A single experimentation q' leads the system from $\omega(q)$ to $\omega(q')$ with positive probability if and only if (regardless of whether $\mu = 1$ or $\mu \in (0,1)$)

$$f(q,q') := \pi(q',(n-1)q'+q) - (\frac{n-1}{nk-1}\pi(q,(n-1)q'+q) + \frac{n(k-1)}{nk-1}\pi(q,nq)) \geq 0. \tag{2.5}$$

The first element in $f(q,q')$ is the profit of the experimenter, which must exceed the average of the profits of the competitors (the first term in brackets) and the non-competitors (the second term in brackets), who all produce q since we are in an absorbing state. Whereas for a transition to be possible when IBM is used, the experimenter's payoff must be better both in relative and in absolute terms, here it is then sufficient that it be better than a weighted average of the profits of the firms in the own market and those in other markets. The weights are determined by n and k. In what follows we write $\lambda := \frac{n-1}{nk-1}$, where we note that $\lambda \in (0,1/2)$ and that $\partial\lambda/\partial k < 0$, $\partial\lambda/\partial n > 0$, $\lim_{n\to\infty}\lambda = 1/k$ and $\lim_{k\to\infty}\lambda = 0$.

If we compare (5) with (3) it is directly seen that $f(q,q') \geq 0$ is less demanding than $D(q,q') \geq 0$ and we can thus expect fewer absorbing states to be stable against single experimentations in the case of IBA. This expectation holds true. As we will show, a single absorbing state, which we will denote by $\omega(q^s)$, remains stable against single experimentations when we consider IBA. This means that q^s in some sense (to be made precise in the following section) is an evolutionary stable strategy. It turns out that q^s also has the property of being able to invade any other absorbing state. These results imply that $\omega(q^s)$ is stochastically stable and that waiting times to arrive at it are as low as they can get. Furthermore, local experimentation tends to drive the system ever closer to $\omega(q^s)$. In what follows we will study the function $f(q,q')$ in order to prove these claims.

2.4.1 The Evolutionary Stable Strategy

If we fix the first argument of $f(q,q')$, the points q' at which $f(q,q') \geq 0$ coincide with the strategies that are capable of invading $\omega(q)$. If there is some q such that

$f(q, q') < 0$ for all $q' \neq q$, then this implies that there is no strategy capable of invading $\omega(q)$. Such an uninvadeable quantity is related to the concept of finite population evolutionary stable strategy by Schaffer (1988) mentioned in section 3.1. An adaptation of this concept to the multi-market context, taking into account average payoffs, renders the following:

Definition 2.3. *The quantity q is an evolutionary stable strategy with respect to average payoffs if and only if $f(q, q') < 0$, for all $q' \in \Gamma$, $q' \neq q$.*

If a strategy q is evolutionary stable according to this definition it follows that a single experimentation will be lead back by the imitation dynamics to $\omega(q)$ with probability one. We now prove the existence and uniqueness of an evolutionary stable strategy and provide the expression that defines it:

Proposition 2.3. *There is a unique evolutionary stable strategy with respect to average payoffs $q^s(k, n)$ defined by $f_2(q^s(k, n), q^s(k, n)) := P'(nq^s(k, n))(1 - \lambda)q^s(k, n) + P(nq^s(k, n)) - C'(q^s(k, n)) = 0$. Furthermore $q^s(k, n) \in (q^c, q^w)$.*[11]

Proof. Let $h_q(q')$ be the function obtained by fixing the first argument of $f(q, q')$ at q. We will show that there is a unique quantity q^s such that $h_{q^s}(q') < 0$ for all $q' \neq q^s$. We do this in 2 steps: (i) We show that there exists a unique quantity q^s such that $h'_{q^s}(q^s) = 0$ and that $q^s \in (q^c, q^w)$. Since this implies $h'_q(q) \neq 0 \; \forall q \neq q^s$ and we know that $f(q, q) = 0 \; \forall q$, the consequence is that for all $q \neq q^s$ there exists some q' such that $h_q(q') > 0$. Therefore, no quantity other than q^s can be evolutionary stable. (ii) We study the function $h_{q^s}(q')$ and confirm that it is increasing to the left of q^s and decreasing to the right, which then guarantees that $h_{q^s}(q') < 0$ for all $q' \neq q^s$.

(i) We differentiate $h_q(q')$ and obtain $h'_q(q') = P'((n-1)q + q')(q' - \lambda q) + P((n-1)q + q') - C'(q')$. Now, let $g(q) := h'_q(q) = P'((nq)(1 - \lambda)q + P(nq) - C'(q)$. We first

[11]In what follows we omit the arguments of $q^s(k, n)$ whenever possible and write q^s, to slim notation.

note that $g(q)$ is strictly decreasing since $g'(q) = nP''((nq)(1-\lambda)q + P'((nq)(1-\lambda) + P'(nq) - C''(q') < 0$. We can also find two points at which $g(q)$ is positive and negative respectively; $g(q^c) = P'((nq^c)q^c + P(nq^c) - C'(q^c) - P'((nq^c)\lambda q^c = -P'((nq^c)\lambda q^c > 0$ and $g(q^w) = P'((nq^w)(1-\lambda)q^w + P(nq^w) - C'(q^w) = P'((nq^w)(1-\lambda)q^w < 0$. By continuity there is thus a unique quantity q^s which lies in (q^c, q^w) such that $g(q^s) = h'_{q^s}(q^s) = 0$ and in fact $g(q) = h'_q(q) > 0$ for $q < q^s$ and $g(q^s) = h'_q(q) < 0$ for $q > q^s$.

(ii) We first note that $h_q(q')$ is concave for $q' > \lambda q$. By differentiating twice we obtain $h''_q(q') = P''((n-1)q + q')(q' - \lambda q) + 2P'((n-1)q + q') - C''(q)$ which is negative as long as $q' \geq \lambda q$.

Next, we observe that $h_q(q')$ is increasing when $q' \leq \lambda q \leq \lambda q^w$ and $(n-1)q + q' < nq^w$. This follows since $P'((n-1)q+q')(q'-\lambda q) > 0$ and $P((n-1)q+q') - C'(q') \geq 0$ (since $q' \leq q \leq nq^w$), which guarantees that $h'_q(q') = P'((n-1)q + q')(q' - \lambda q) + P((n-1)q + q') - C'(q') > 0$.

Since $q^s < q^w$, we therefore have that $h_{q^s}(q') < 0$ for all $q' \neq q^s$ □

The intuition behind this result is that there are experimentations from $\omega(q^s)$ that increase profits in absolute terms (toward the Cournot quantity) and experimentations generating advantages in relative terms (toward the Walrasian outcome). However, q^s is precisely such that the absolute advantages obtained by experimenting to lower quantities are outweighed by the disadvantages in terms of relative payoffs, and vice-versa for experimentations to higher quantities. It turns out that in the present setting such a quantity always exists and is unique. Figure 2.1 illustrates $f(q, q')$ in a linear setting, fixing the the first argument:

As indicated by the preceding plot, q^s can also be obtained as a solution to a maximization problem. This follows as a direct corollary to the Proposition 2.3:

Corollary 16. *The evolutionary stable strategy q^s is the unique solution to $q \in \arg\max_{q'} f(q, q')$.*

Figure 2.1: Invadeable Quantities in a Linear Setting

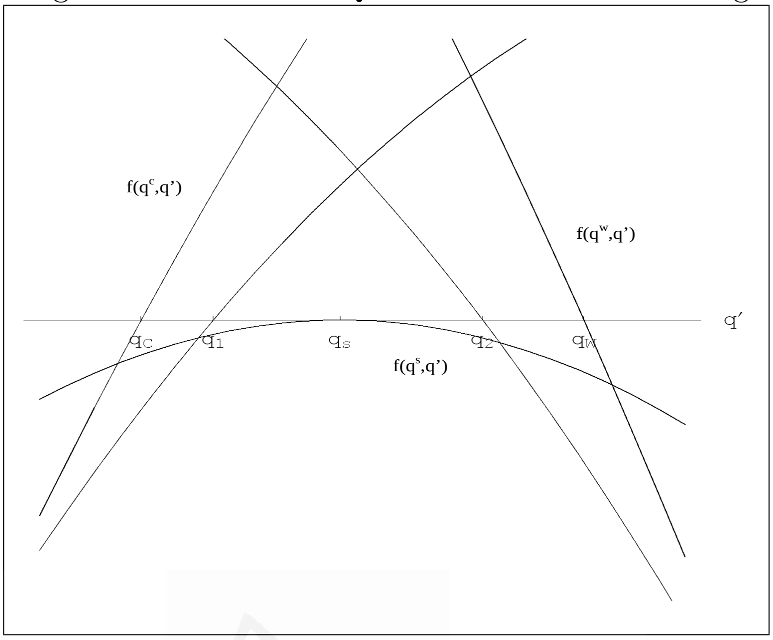

Proof. We take first order necessary conditions and evaluate at $q' = q$, obtaining $P'(nq)(1-\lambda)q + P(nq) - C'(q)$. This corresponds to the function $g(q)$ in the proof of Proposition 2.3, where it is noted that $g(q) = 0$ if and only if $q = q^s$. The sufficient condition is obtained in part (ii) of the proof, where it is seen that $f(q^s, q')$ is increasing for $q' < q^s$ and decreasing for $q' > q^s$. $\square$

2.4.2 The Global Invader

If instead of fixing the first argument of $f(q, q')$ we fix its second argument, then the points q at which $f(q, q') > 0$ gives the states $\omega(q)$ that can be invaded by single experimentations to q'. These points represent the states from which an experimentation to q' gives the experimenter the highest average payoff. If there is some q' such that $f(q, q') > 0$ for all $q \neq q'$, then we say that q' is a global invader. We now proceed to show that q^s is not only immune to invasion, but that it is indeed also a global invader, so a single experimentation to q^s from any state $\omega(q)$ moves the system with positive probability to $\omega(q^s)$.

Proposition 2.4. *q^s is the unique global invader.*

Proof. Uniqueness is a consequence of the preceding Proposition. Since $f(q^s, q') < 0$ $\forall q' \neq q^s$ no strategy invades q^s and therefore no strategy except q^s can be a global invader. We thus proceed to prove that q^s is a global invader. We will do this in 3 steps. Let $f_{q'}(q)$ be the function obtained by fixing the invading strategy q' in $f(q, q')$. We will show that (i) $f'_{q^s}(q^s) = 0$; (ii) $f'_{q^s}(q) > 0$, $\forall q > q^s$; (iii) $f'_{q^s}(q) > 0$ $\forall q < q^s$. Since we know that $f(q, q) = 0$ $\forall q \in \Gamma$, it is then assured that $f_{q^s}(q) > 0$ $\forall q \in \Gamma$.

(i) Write $f'_{q'}(q) = (n-1)P'((n-1)q + q')(q' - \lambda q) - \lambda P((n-1)q + q') - (1-\lambda)nP'(nq)q - (1-\lambda)P(nq) + C'(q')$, which means that $f'_{q'}(q') = -P'((nq')q'(1-\lambda) - P(nq') + C'(q') = -g(q') = 0$ only when $q' = q^s$.

(ii) Consider $f''_{q'}(q) = (n-1)^2 P''((n-1)q+q')(q'-\lambda q) - 2\lambda(n-1)P'((n-1)q+q') - (1-\lambda)n^2 P''(nq)q - 2(1-\lambda)nP'(nq) + C''(q')$. All terms except the first one are positive and the first one is positive for $\lambda q > q'$. Hence $f''_{q'}(q) > 0$ and $f_{q'}(q)$ is therefore convex for $q > q'/\lambda$.

We now show that $f'_{q'}(q) > 0$, for $q \in (q', q'/\lambda)$, $q' \geq q_s$. Consider the cross derivative $f_{12}(q, q')$. This can be thought of as the change in $f'_{q'}(\hat{q})$ for a fixed $q = \hat{q}$ as we vary q'. We compute $f_{12}(q, q') = (n-1)P''((n-1)q + q')(q' - \lambda q) + (n - 1 - \lambda)P'((n-1)q + q')$ and note that this is negative as long as $q \leq q'/\lambda$. We know that $f_1(q, q) > 0$ for $q > q_s$. Then, take any $q' \geq q_s$ and $q \in (q', q'/\ \lambda)$. Since $f_{12}(q, q')$ is negative in $(q', q'/\ \lambda)$ we must have that $f_1(q, q') > f_1(q, q) > 0$. Hence, $f'_{q'}(q) = f_1(q, q') > 0$ when $q' \geq q_s$ and $q \in (q', q'/\lambda)$.

Since $f'_{q^s}(q) > 0$, for $q \in (q_s, q_s/\lambda)$ and it is convex when $q > q_s/\lambda$, it follows that $f'_{q^s}(q) > 0$, $\forall q > q^s$.

(iii) Again we use the cross derivative $f_{12}(q, q')$. Take any $q < q' \leq q_s$. We know that $f_1(q, q) < 0$, for $q < q'$. Since $f_{12}(q, q') = (n-1)P''((n-1)q + q')(q' - \lambda q) + (n - 1 - \lambda)P'((n-1)q + q')$ is negative for all $q \leq q'/\lambda$, we have that $f_1(q, q)$ is decreasing in its second argument. Hence $f_1(q,q') < f_1(q,q) < 0$ for all $q < q' \leq q_s$, which implies $f'_{q^s}(q) = f_1(q,q^s) < 0$, $\forall q < q^s$. $\square$

An observation of the plots in Figure 2.1 indicates that quantities $q < q^s$ are invaded by $q' > q$, whereas $q > q^s$ are invaded by $q' < q$. This turns out to be true beyond the linear case. There is a tendency for the system to move in small steps towards q^s, since any strategy is invaded by quantities in the direction of q^s but not in the opposite direction. This is a consequence of the proofs of Proposition 2.3 and 2.4 and is therefore here presented as a corollary:

Corollary 17. *(i) Any $\omega(q)$, $q < q^s$ is invaded by $(q,q^s]$, but not by $[0,q)$. (ii) Any $\omega(q)$, $q > q^s$ is invaded by $[q^s,q)$, but not by $(q,v\delta]$.*

Proof. (i) We know that $f_{q'}(q') = 0\ \forall q'$. In part (iii) of the proof of Proposition 2.4 we showed that $f'_{q'}(q) < 0$ for all $q \leq q' < q_s$ (the non strict inequality $q \leq q'$ follows from $f'_{q'}(q') < 0\ \forall q' < q^s$). Therefore $f_{q'}(q) > 0\ \forall q,q'$ such that $q \leq q' \leq q_s$. This means that any $q' \leq q^s$ invade all the strategies to its left and hence $\omega(q)$, $q < q^s$ are invaded by $(q,q^s]$. Next, we know from part (ii) of the proof of Proposition 2.3 that $h_q(q')$ is concave for $q' \geq \lambda q$ and that $h'_q(q') > 0$ when $q' \leq \lambda q$. Since $h_q(q) = 0$ and $h'_q(q) > 0$ we can thus be sure that $h_q(q') < 0$, $\forall q' < q$, which means that q is not invaded by any strategy $[0,q)$.

(ii) Part (ii) of the proof of Proposition 2.4 directly gives $f'_{q'}(q) = f_1(q,q') > 0$ $\forall q',q$ such that $q' \geq q > q^s$. This means that any $q' \geq q^s$ invades all strategies to its right and hence $\omega(q)$, $q < q^s$ are invaded by $[q^s,q)$. Next, $h_q(q) = 0$, $h'_q(q) < 0$, $\forall q > q^s$ together with the concavity of $h_q(q')$ for $q' \geq \lambda q$ implies $h_q(q') < 0\ \forall q,q'$ such that $q' \geq q > q^s$ and q is therefore not invaded by $(q,v\delta]$. $\square$

Figure 2.2 illustrates the function $f(q,q')$ in a linear setting holding the second argument fixed, for some different values of q'.

Figure 2.2: Invading Quantities in a Linear Setting

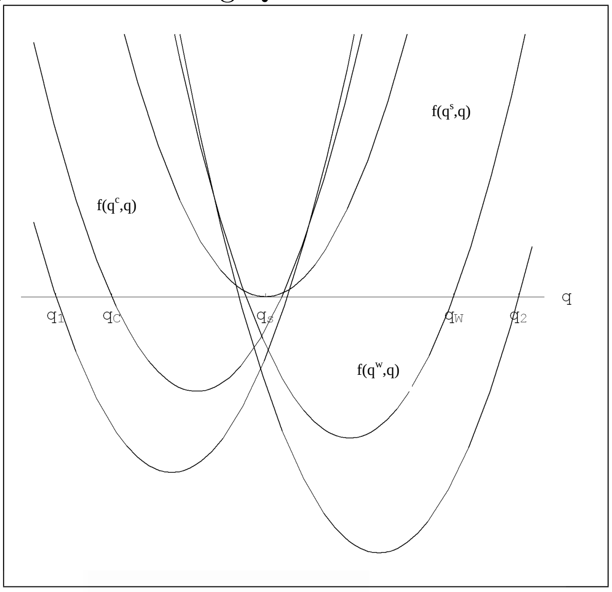

2.4.3 Stochastic stability

Given what has been shown so far, it should come as no surprise that $\omega(q^s)$ is stochastically stable. From any absorbing state $\omega(q)$ it is possible to reach $\omega(q^s)$ by means of a single experimentation, but a single experimentation is never enough to leave this state.

Proposition 2.5. *When firms use IBA $\omega(q^s)$ is the unique stochastically stable state and the expected waiting time until convergence to $\omega(q^s)$ is in the order of magnitude of ϵ^{-1}.*

Proof. We use a radius-coradius argument (Appendix A). By Proposition 2.3, $R(\omega(q^s)) > 1$ and by Proposition 2.4, $Cr(\omega(q^s)) = 1$. Hence $R(\omega(q^s)) > Cr(\omega(q^s))$ and $\omega(q^s)$ is stochastically stable. The expected waiting time until convergence to $\omega(q^s)$ is in the order of magnitude of $\epsilon^{-Cr(\omega(q^s))} = \epsilon^{-1}$. □

In the long run the process will thus spend almost all of its time in $\omega(q^s)$. The expected waiting time is as low as it can get in this kind of models, which means that the speed of convergence is relatively high. As a consequence $\omega(q^s)$ is a reasonable prediction also of what we can expect to see in the medium run.

This result should be contrasted with Apesteguia et. al. (2007), who obtain that when firms use IBA, $\omega(q^c)$ is the unique stochastically stable quantity. However this result depends on the assumption of random remixing of the markets from one period to another, as well as the specific linear setting analyzed, in which only five possible quantities can be chosen. By relaxing these assumptions we obtain a different prediction. The quantity q^s also depends on the number of markets, a feature that does not arise in Apesteguia et. al. (2007).

2.4.4 Stochastic Stability Under Local Experimentation

We can use the property that experimentations in the direction of q^s are imitated but not experimentations in the opposite direction (Corollary 17), to prove that $\omega(q^s)$ is stochastically stable under local experimentation. By local experimentation we mean that firms only experiment with small changes in production. Experimentation is likely to be local if firms are conservative and reluctant to change output too abruptly. To capture the idea of local experimentation we therefore assume that a firm ij picked to experiment chooses quantity according to some distribution over $\{[q_{ij} - \alpha, q_{ij} + \alpha] \cap \Gamma\}$, where $\alpha > \delta$, with full support.

In the case of global experimentation it was argued that since a transition from any state x to any state y occurs with positive probability in the perturbed process, this is guaranteed to be ergodic. With local experimentation we cannot use this argument. Nevertheless, with local experimentation any state y is reachable from any state x by a sequence of experimentations in m steps, and this is also sufficient for ergodicity. Hence, the process remains ergodic, which means that we can apply the standard results in the literature to characterize the set of stochastically stable states[12]. We then obtain:

[12] As mentioned by Vega-Redondo (2003) p.p. 477, these results continue to hold as long as the perturbed process is ergodic, even if there are some transitions that occur with probability

Proposition 2.6. *The state $\omega(q^s)$ is the unique stochastically stable state when experimentation is local.*

Proof. We use a radius - modified coradius argument (see Appendix A) to prove this result.

In Corollary 17 it was shown that experimentations in the direction of q^s are always imitated. If we start from any $\omega(q)$, $q < q^s$, a state $\omega(q') = \omega(q+\alpha)$ can thus be reached by one experimentation. From $\omega(q')$ a state $\omega(q'+\alpha)$ can be reached by one experimentation. In this way, $\omega(q^s)$ can be reached by a series of successive transitions between absorbing states, each requiring a single experimentation. Since each absorbing state reached in this sequence of transitions is left with a single experimentation, it has radius equal to one. We therefore obtain $Cr^*(\omega(q^s)) = 1$. The same argument can be repeated if $q > q^s$. At the same time, by Proposition 2.4 $\omega(q^s)$ cannot be left with single experimentations, which means that $R(\omega(q^s)) > 1$. Consequently, $R(\omega(q^s)) > Cr^*(\omega(q^s))$ and $\omega(q^s)$ is therefore the unique stochastically stable state when experimentation is local. □

This result shows that the stochastic stability of $\omega(q^s)$ does not depend on q^s being a global invader. In fact, if firms are prudent and experiment only locally, the system will tend to move in small steps towards q^s.

2.4.5 Comparative Statics

Since we have an expression that defines q^s implicitly we can perform a comparative statics exercise to evaluate how this quantity depends on the number of firms in each markets and the number of markets.

Proposition 2.7. *$q^s(k,n)$ decreases in n and k, $\lim_{k\to\infty} q^s(k,n) = q^c$.*

zero. I.e. the sufficient condition for ergodicity that a transition between any two states ocurrs with positive probability is not necessary for the standard results in the literature on stochastic stability to hold.

Proof. We differentiate $q^s(k,n)$ implicitly with respect to k using the expression $P'(nq^s)(1-\lambda)q^s + P(nq^s) - C'(q^s) = 0$ and obtain $\frac{\partial q^s}{\partial k} = -\frac{P'(nq^s)q^s}{g'(q^s)}\frac{n(n-1)}{(nk-1)^2} < 0$. Next, note that $\frac{\partial \lambda}{\partial n} = \frac{k-1}{nk-1} \in (0,1)$, so by implicit differentiation we obtain $\frac{\partial q^s}{\partial n} = -\frac{P''(nq^s)(1-\lambda)(q^s)^2 - P'(nq^s)q^s\frac{\partial \lambda}{\partial n} + P'(nq^s)}{g'(q^s)} = -\frac{P''(nq^s)(1-\lambda)(q^s)^2 + P'(nq^s)q^s(1-\frac{\partial \lambda}{\partial n})}{g'(q^s)}$. Both the denominator and the numerator of this expression are negative so $\frac{\partial q^s}{\partial n} < 0$. Further, we know $\lim_{k\to\infty}\lambda = 0$ which means $\lim\limits_{k\to 0} P'((nq^s)(1-\lambda)q^s + P(nq^s) - C'(q^s) = P'((nq^s)q^s + P(nq^s) - C'(q^s)$ and if we put this expression equal to zero we obtain $\lim\limits_{k\to\infty} q^s(k,n) = q^c(n)$. □

When considering these comparative statics it should be noted that the quantities q^c and q^w do not change with k. As the number of markets changes, q^s thus moves in the fixed interval (q^c, q^w), which means that the effect on competition can be evaluated straightforwardly. In standard Cournot oligopoly models with optimizing firms, the number of markets has no effect at all. Firms will not consider what happens in other markets since this is irrelevant for their optimization problem. Here, more markets make relative payoff considerations less important, thereby giving weight to the Cournot outcome. An implication of the results derived here is that if an additional local market for a certain good appears within a city (keeping the number of firms per market fixed), competition in each market will actually decrease, even if the total number of firms operating within the city increases. This effect, which is in stark contrast to traditional predictions, could have interesting consequences for competition policy.

The prediction provided by the comparative statics exercise is particularly suitable for tests in the laboratory, since it provides a direct link between a parameter that is very easy to vary and the outcome. This is a possibility was not exploited in the experiments carried out by Apesteguia et. al. (2007) (given that their model does not generate such a result). An oligopoly experiment can be designed in which individuals are grouped into different markets and have information about what people in the same and in other markets choose and earn. The number of markets is varied across sessions. If the outcome is unaffected by the number of markets, this indicates that people do not imitate across markets according to IBA.

When it comes to the parameter n, the analysis becomes a bit more complicated. First, both q^c and q^w are functions of n. With assumptions 1-3 q^c and q^w can be shown to decrease in n. On the other hand, the total competitive output nq^w is constant in n, and nq^c approaches nq^w as n becomes large (see for example Shapiro (1989)). Since $q^s \in (q^c, q^w)$ we can thus be certain that the total output of each market when firms use IBA approaches the competitive one as the number of firms per market becomes large. This will occur through two different mechanisms. On one hand, given any state $\omega(q)$ the parameter n will determine the effect on the own and the competitors profits of a given experimentation to q', via $P(Q)$. On the other hand, n will affect the weight λ in the computation of the average payoff of non-experimenting firms. We can note that $\frac{\partial q^s}{\partial \lambda}\frac{\partial \lambda}{\partial n} > 0$. This expression in a sense "isolates" the effect of n on q^s that goes through λ. The effect of n on λ thus makes q^s creep closer to q^w. Hence the net effect of n is to increase the weight of the direct competitors in the average payoff computation, thus making relative payoff comparisons more important.

2.4.6 An Alternative Interpretation of q^s

Corollary 16 can be used to find settings in which firms display optimizing behavior and that generate the same prediction as IBA. One possibility is to consider a single market of firms that care about their profit in absolute terms as well as about outdoing the competition. Consider a single market and let the utility of firms be:

$$U(q_i, q_{-i}) = (1-\lambda)\pi(q_i, Q) + \lambda(\pi(q_i, Q) - K(\pi(q_{-i}, Q))), \tag{2.6}$$

where $K(\pi(q_{-i}, Q))$ is a convex combination of the profits of all firms but i. A firm's utility is then a weighted average of its profit in absolute terms and its advantage over other firms. The parameter λ captures the extent to which firms prefer absolute versus relative payoffs. If $\lambda = 0$ we obtain a standard Cournot game, whereas if $\lambda = 1$ we obtain a game of relative payoff maximization. Define a symmetric Nash equilibrium as a quantity q such that $q \in \arg\max_{q'}\{U(q', (n-1)q + q)\}$. We then have

Proposition 2.8. *$q^s(k,n)$ such that $\lambda = \frac{n-1}{nk-1}$ is the unique symmetric Nash equilibrium of the game in which firms' utility is captured by $U(q_i, q_{-i})$ and it is also a strict Nash equilibrium.*

Proof. We write $U(q', (n-1)q + q') = (1-\lambda)\pi(q', (n-1)q + q') + \lambda(\pi(q', (n-1)q + q') - \pi(q, (n-1)q + q')) =$

$\pi(q', (n-1)q + q') - \lambda\pi(q, (n-1)q + q') = f(q, q') + (1-\lambda)\pi(q, nq)$. For a symmetric Nash equilibrium we then need $q \in \arg\max_{q'} f(q, q') + (1-\lambda)\pi(q, nq)$. But the problem $\max_{q'} f(q, q') + (1-\lambda)\pi(q, nq)$ is equivalent to $\max_{q'} f(q, q')$, since $(1-\lambda)\pi(q, nq)$ is just a constant term in this maximization problem. We then need to solve $q \in \arg\max_{q'} f(q, q')$ and from 16 we know that q^s is the unique solution to this problem. Further, since $f(q^s, q') < 0$ for $q' \neq q^s$ the equilibrium is strict. □

Proposition 2.6 relates the prediction obtained when firms use IBA to optimizing behavior. Imitating the best average in multiple markets corresponds to a prediction of a model in which firms best respond and are concerned about outdoing the competition, but at the same time have an interest in their absolute payoffs.

2.5 Alternative Informational Settings

2.5.1 Imitate Only if the Aggregates are Sufficiently Close

Intuitively, imitation makes sense if the circumstances of the sampled firm are sufficiently similar to those of the imitator. Even though in this paper we consider markets that are ex-ante identical, they may differ ex-post with respect to the aggregate quantity produced. If the aggregates are very different, the profit obtained by a certain strategy in another market may be felt to be a poor predictor of what is a successful strategy in the own market. It should be noted, however, that the profit obtained in a market with a different aggregate is not *necessarily* a worse predictor than that obtained by a firm in the own market. The reason is that when imitating a firm in the own market, the ex-post aggregate will differ from that before imitation took place. For example, following an experimentation from a state $\omega(q)$ that occurs in another market than the own, which implies imitating although the aggregates are different, is always more informative about what profits will be obtained than if the experimentation had occurred in the own market.

In this section we explore the implications of assuming that firms are somewhat cautious with respect to imitating firms in other markets, and are willing to imitate these only if the difference in the aggregates is below a certain level. This corresponds to imitation that is local, rather than global, in the sense that not all firms are necessarily sampled. However, localness is endogenous, since it depends on the actual quantities produced. We show that the outcome tends to become more competitive in this case, the reason being that markets will evolve more independently, reinforcing the relative payoffs effect. We provide a full characterization of the outcome when IBM is used, and a partial characterization when IBA is used, in which case more analytical difficulties arise.

The Model

We assume that each time a firm ij is picked to imitate it samples firms in markets j' such that $|Q_j - Q_{j'}| \leq \bar{Q}$. $\bar{Q}$ parameterizes local imitation and we obtain a model in which firms gradually become more willing to imitate across markets as $\bar{Q}$ increases. For large values of $\bar{Q}$ we obtain a model similar to the multimarket model analyzed in the preceding sections, and for small values of $\bar{Q}$ markets evolve more independently, resembling Vega-Redondo (1997). For intermediate values we obtain a model between these two extremes.

We will work with a linear model, with $P(Q) = a - bQ$ and $C(q) = cq$, in order to obtain more clear-cut results. Similar results as those presented here can be obtained without this last assumption. We prefer to present the results for a linear model since the results are clearer in this way (even in the linear case, the proof of stochastic stability is a bit lengthy) and having more general demand and costs does not provide any additional intuition.

Results

We start by noting that the set of absorbing states will be slightly different in this setting. Now, firms in different market need not necessarily produce the same quantity in absorbing states. Apart from the "monomorphic" states $\{\omega(q) : q \in \Gamma\}$, "polymorphic" states in which all firms in a given market produce the same quantity, but the difference between quantities produced in different markets is sufficiently large, will now also be absorbing. We denote by a vector $\mathbf{q} = (q_1, q_2, ..., q_k)$ a state in which q_j is produced by each firm in market $j \in K$. Let $\Omega_P := \{\mathbf{q} \in \Gamma^k : \forall j, j' \in K,$ $q_j = q_{j'}$or $|Q_j - Q_{j'}| > \bar{Q}\}\backslash\Omega_M$, (where subscript P stands for "polymorphic"). As before, Ω denotes the set of absorbing states. Then:

Lemma 2.2. *The set of absorbing states Ω of the unperturbed process is $\Omega_M \cup \Omega_P$.*

Proof. (i) Any state in Ω_M is obviously an absorbing state since there is nothing to imitate. Any state in Ω_P is also absorbing, since firms that are close enough produce the same quantity, and the remaining firms are too far away to be imitated.

(ii) It remains to show that any state in which there is some $ij, i'j' \in N \times K$ such that $|q_{ij} - q_{i'j'}| \leq \bar{Q}$, $q_{ij} \neq q_{i'j'}$ is not absorbing. Consider first a state in $\{\mathbf{q} \in\Gamma^k\}\backslash\Omega_P$ in which x different quantities are produced. Take some j and j' such that $|Q_j - Q_{j'}| \leq \bar{Q}$. Then there is always a positive probability that either all firms in all markets producing q_j imitate $q_{j'}$ or vice versa, in which case we reach a state in $\{\mathbf{q} \in\Gamma^k\}\backslash\Omega_P$ in which $x-1$ quantities are produced. By iteration of this argument we reach a state in $\Omega_M \cup \Omega_P$. Next, consider a state in $\Gamma^{kn}\backslash\{\mathbf{q} \in\Gamma^k\}$. From such a state there is always a positive probability that all firms in a market imitate the same quantity, in which case we reach a state in $\{\mathbf{q} \in\Gamma^k\}$. Hence, there is always a positive probability of going from a state in which there is some $ij, i'j' \in N \times K$

such that $|Q_{ij} - Q_{i'j'}| \leq \bar{Q}$, $q_{ij} \neq q_{i'j'}$ to a state in $\Omega_M \cup \Omega_P$ and such states are therefore not absorbing. $\square$

Next, we characterize the set of stochastically stable states of the perturbed process:

Proposition 2.9. *The set of stochastically stable states is* $\{\omega(q) : q \in [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$.

Proof. We will use the shorthand "a series of x-ET" for " a series of x-experimentation transitions", which refer to a sequence of transitions between absorbing states, each of which requires x experimentations. "pp" abbreviates "positive probability".

Step 1: Any state in Ω_P can be left for a state in Ω_M through a series of 1-ET.

Consider a state in Ω_P such that $||\{(j,j') \in K \times K : j \neq j' \text{ and } Q_j \neq Q_{j'}\}|| = x > 0$. We will show that a series of 1-ET reduces $||\{(j,j') \in K \times K : j \neq j' \text{ and } Q_j \neq Q_{j'}\}||$. Take markets j and $j' : Q_{j'} \in \min_{j'' \in K}\{Q_{j''} : Q_{j''} > Q_j\} \neq \varnothing$. Let a firm ij experiment to $q' : (n-1)q_{ij} + q' = Q_{j'} - \bar{Q}$. Then markets j and $j'' : Q_{j''} = Q_{j'}$ will be within comparison reach and all firms in these markets thus imitate the same quantity with pp. If $j'' : Q_{j''} = Q_{j'}$ imitate j' we immediately reach a state in which $||\{(j,j') \in K \times K : j \neq j' \text{ and } Q_j \neq Q_{j'}\}|| < x$ with pp. If j' imitate $j'' : Q_{j''} = Q_{j'}$, repeat the same procedure for all markets producing Q_j. We again reach a state in $||\{(j,j') \in K \times K : j \neq j' \text{ and } Q_j \neq Q_{j'}\}|| < x$. By iteration of this argument we reach a state in which $x = 0$.

We have thus showed $\Omega_P \xrightarrow{1} ... \xrightarrow{1} \Omega_M$.

Step 2: Any state $\omega(q)$ such that $q' \in \Gamma \backslash [q^c, q^w]$ can be left for an absorbing state in which $q \in [q^c, q^w]$ through a series of 1-ET.

(i) Consider a state $\omega(q)$ such that $q > q^w$. Let some firm experiment to $\max\{q^w, q - \bar{Q}\}$. This experimentation is imitated by all firms with pp. Repeat until $\omega(q^w)$ is reached.

(ii) Consider a state $\omega(q)$ such that $q < q^c$. Let some firm experiment to $\min\{q^c, q + \bar{Q}\}$. This experimentation is imitated with pp. Repeat until $\omega(q^c)$ is reached.

Let $\Omega_1 := \{\omega(q) : q \notin [q^c, q^w]\}$. We have then shown $\Omega_1 \xrightarrow{1} ... \xrightarrow{1} \Omega_M \backslash \Omega_1$.

Step 3: Any state in $\{\omega(q) \in \Omega_M \backslash \Omega_1 : q \notin [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$ can be left for a state in $\Omega_2 := \{\omega(q) : q \in [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$ through a series of 1-ET. No 1-ET from Ω_2 is possible.

(i) Consider a state in $\{\omega(q) \in \Omega_M \backslash \Omega_1 : q \notin [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$. We will show that there are experimentations q' that are profitable in relative terms, i.e. $d(q, q') := \pi(q', (n-1)q + q') - \pi(q, (n-1)q + q') \geq 0$ and that bring the experimenting market out of comparison reach of the remaining markets. This means that absolute payoff considerations become irrelevant and that the experimentation is followed with pp. Formally, it is required that $d(q, q') = (a - b((n-1)q + q') - (cq' - cq) \geq 0$ and $q' - q > \bar{Q}$. $d(q, q')$ is a concave function with $d(q, q) = 0$, $d(q, \varphi(q)) = 0$ at some point $\varphi(q) > q^w$ for $q < q^w$ and $d(q, q') > 0$, for $q' \in (q, \varphi(q))$.

Hence, for any quantity $q < q^w$ there is an upper bound $\varphi(q)$ to experimentations such that $d(q, q') \geq 0$. We deduce $\varphi(q) = \frac{a-c}{b} + (1-n)q$. The mentioned experimentations are possible as long as $q' \leq \frac{a-c}{b} + (1-n)q$ and $q' > \bar{Q} + q$. Both inequalities hold for $q < \underline{q} = q^w - \bar{Q}/n$. If $\underline{q} < q^c$ then the mentioned transition is not possible for any state in $\Omega_M \backslash \Omega_1$. This transition is thus possible from any state $\{\omega(q) \in \Omega_M \backslash \Omega_1 : q \notin [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$, but not from states in Ω_2.

(ii) Consider a state in $\{\omega(q) \in \Omega_M \backslash \Omega_1 : q \notin [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$. We will show how a series of 1-ET leads the system into $\omega(q^w)$. Let a firm ij experiment to $\varphi(q)$. The firms in j will then imitate with pp. and we reach a new absorbing state (in Ω_P). Next, let a firm $i'j'$ experiment to $\varphi(q)$. j' will not be in comparison reach of j before imitation takes place, since this requires $(n-1)q + q' + \bar{Q} \geq nq' \leftrightarrow q' - q < \bar{Q}/(n-1)$ and we required $q' - q > \bar{Q}$. All firms in j' then imitate q' with pp. By proceeding in this way we reach the state $\omega(\varphi(q))$. From this state let a firm ij experiment to $\max\{q^w, \varphi(q) - \bar{Q}\}$. Since $D(\varphi(q), \max\{q^w, \varphi(q) - \bar{Q}\}) > 0$ all firms imitate $\max\{q^w, \varphi(q) - \bar{Q}\}$ with pp. and we thus reach $\omega(\max\{q^w, \varphi(q) - \bar{Q}\})$. Proceed like this until $\omega(q^w)$ is reached. Hence, starting from a state in $\{\omega(q) \in \Omega_M \backslash \Omega_1 : q \notin [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$ we reach the state $\omega(q^w)$ through a series of 1-ET.

(iii) Since in any state in Ω_2 there are no experimentations such that $d(q, q') > 0$ and $q' - q > \bar{Q}$, a 1-ET can only occur if $D(q, q') > 0$. But we know from Alós-Ferrer (2004) Lemma A6 that $D(q, q') < 0$ for all $q \in [q^c, q^w]$, $q' \neq q$, so no one-experimentation transition out of Ω_2 is possible.

Step 4: All states in Ω_2 can be connected through a series of 1- and 2-experimentation transitions.

(i) First, a 2-ET from a state $\omega(q) \in \Omega_2$ to some state $\omega(q'') \in \Omega_2$, $q < q''$ is always possible. Consider a state in Ω_2 and two simultaneous experimentations $q' = 0$ and $q'' > q$. Then q'' is imitated with pp. iff $e(q, q'') := (a - b((n-2)q + q'')(q'' - q) - (cq'' - cq) \geq 0$. For $q < q^w$ the function $e(q, q'')$ is concave with $e(q, q) = e(q, nq^w - (n-2)q) = 0$ and $e(q, q'') > 0$ for all $q'' \in [q, nq^w - (n-2)q\}]$. Any $q'' \in (q, \max\{nq^w - (n-2)q, 2q + \bar{Q}\}]$ is thus imitated with pp. by all firms, where the second entry of the maximum is the requirement that the experimenting market does not exit comparison reach. Hence, we have that $\omega(q) \xrightarrow{2} \omega(q'')$ is possible for any q, q'' such that $\max\{q^c, q^w - \frac{\bar{Q}}{n}] \leq q < q'' < \max\{nq^w - (n-2)q, 2q + \bar{Q}\}$.

(ii) However, there are no corresponding downward 2-ET (Alós-Ferrer (2004) Lemma B1). Instead, we will show that a 2-ET from $\omega(q^w)$ to some quantity $q'' > q^w$, from which a 1-ET to any quantity in $[q^w - \frac{\bar{Q}}{n}, q^w]$ is possible.

Since $e(q^w, 2q^w) = 0$, from $\omega(q^w)$ two experimentations $q' = 0$ and $q'' = 2q^w$ means q'' is followed by all firms (this experimentation will not bring the experimenting market out of comparison reach since the aggregate is held constant) with pp. Hence $\omega(q^w) \xrightarrow{2} \omega(\min\{2q^w, \underline{q} + \bar{Q}\})$ (with $\underline{q} := \max\{q^c, q^w - \frac{\bar{Q}}{n}])$ is possible. Next, from $\omega(\min\{2q^w, \underline{q} + \bar{Q}\})$ any experimentation to quantities $q' \in [q^c, \min\{2q^w, \underline{q} + \bar{Q}\})$ improves absolute payoffs (since they are better responses) and straightforward calculus shows that $d(2q^w, q') > 0$ for all $q' \in [q^c, 2q^w)$, which means that an experimentation to $q' \in [q^c, 2q^w)$ is better also in relative terms and is therefore imitated with pp. Correspondingly, $d(\underline{q} + \bar{Q}, q') > 0$ for all $q' \in [\underline{q}, \underline{q} + \bar{Q})$. This means that $\omega(\min\{2q^w, \underline{q} + \bar{Q}\}) \xrightarrow{1} \omega(q)$ is possible for any $q \in [\underline{q}, q^w]$ (we consider

$\min\{2q^w, \underline{q} + \bar{Q}\}$ to be sure that the experimentation back to q does not bring the market out of comparison reach).

Step 5: The set of stochastically stable states is Ω_2.

We will use a combination of radius - modified coradius and tree surgery arguments (see Appendix A) to prove stochastic stability.

(i) First, note that $R(\Omega_2) = 2$. We have shown in step 1-4 that $\Omega_P \xrightarrow{1} ... \xrightarrow{1} \Omega_M \xrightarrow{1} ... \xrightarrow{1} \Omega_M \backslash \Omega_1 \xrightarrow{1} ... \xrightarrow{1} \Omega_2$ is possible. Then $Cr^*(\Omega_2) = 1$. Since $R(\Omega_2) > Cr^*(\Omega_2)$, the set of stochastically stable states is contained in Ω_2.

(ii) We now construct the $\omega(q^w)$-tree, which shows that $\omega(q^w)$ is stochastically stable. We already noted $\Omega_P \xrightarrow{1} ... \xrightarrow{1} \Omega_M \xrightarrow{1} ... \xrightarrow{1} \Omega_M \backslash \Omega_1 \xrightarrow{1} ... \xrightarrow{1} \Omega_2$ and that $\Omega_2 \xrightarrow{2} ... \xrightarrow{2} \omega(q^w)$ (step 4). This means that in the $\omega(q^w)$-tree, all states $\Omega \backslash \Omega_2$ will have exiting arrows of cost 1, while states in $\Omega_2 \backslash \omega(q^w)$ will have exiting arrows of cost 2. Since states in Ω_2 cannot be left with less than 2 experimentations, $\omega(q^w)$ is therefore stochastically stable and the $\omega(q^w)$ tree is $\Omega_P \xrightarrow{1} ... \xrightarrow{1} \Omega_M \xrightarrow{1} ... \xrightarrow{1} \Omega_M \backslash \Omega_1 \xrightarrow{1} ... \xrightarrow{1} \Omega_2 \xrightarrow{2} ... \xrightarrow{2} \omega(q^w)$.

(iii) We now show that any state in Ω_2 has the same stochastic potential as $\omega(q^w)$. Consider a state $\omega(q) \in \Omega_2 \backslash \omega(q^w)$. We know $\{\omega(q') \in \Omega_2 : q' < q\} \xrightarrow{2} ... \xrightarrow{2} \omega(q)$ and at the same time $\{\omega(q') \in \Omega_2 : q' > q\} \xrightarrow{2} ... \xrightarrow{2} \omega(q^w)$ (step 4). Next, $\omega(q^w) \xrightarrow{2} \omega(\min\{2q^w, \underline{q} + \bar{Q}\}) \xrightarrow{1} \omega(q)$. Hence, to obtain the $\omega(q)$-tree we redirect the arrow exiting $\omega(\min\{2q^w, \underline{q} + \bar{Q}\})$ in the $\omega(q^w)$-tree to $\omega(q)$, which will not affect the total cost of the tree. Then we cut the arrow exiting $\omega(q)$, which had cost 2, and add an arrow from $\omega(q^w)$ to $\omega(\min\{2q^w, \underline{q} + \bar{Q}\})$ at cost 2. Neither this will affect the total cost of the tree. Next we make sure there are arrows $\{\omega(q') \in \Omega_2 : q' < q\} \xrightarrow{2} ... \xrightarrow{2} \omega(q)$ (so that none of them "jumps" q directly to $\omega(q^w)$). We thus have a $\omega(q)$-tree with the same stochastic potential as the $\omega(q^w)$-tree and all states in $\{\omega(q) : q \in [\max\{q^c, q^w - \frac{\bar{Q}}{n}], q^w]\}$ are therefore stochastically stable. Together with (i) this proves the result. $\square$

As seen, the proof proceeds in various steps, in which we analyze stability properties of different sets of states. Polymorphic states are inherently instable since large enough experimentations always brings two separated markets together. This might seem like a fragile result, since it appears to depend on firms being willing to experiment to very large and probably unprofitable quantities, but in fact it can be shown that in the linear setting, this result cannot be avoided by for example imposing a capacity constraint[13]. It thus seems quite hard for polymorphic states to survive with local imitation in the sense considered here. States in which quantities not in $[q^c, q^w]$ are produced are also easily destabilized, the reason being (as in the section 3.1) that there are experimentations that improve both relative and absolute payoffs. Further, there is a lower bound $\underline{q} \in [q^c, q^w)$ such that all states in $\{\omega(q) : q < \underline{q}\}$ can be destabilized by an upward experimentation. The intuition is that even if such an experimentation lowers profits in absolute terms, if it is sufficiently large it will bring the experimenting market out of comparison reach from the other markets producing the same quantity. Absolute payoff considerations then become irrelevant and the experimenting market evolves independently.

The set of stochastically stable states depends on $\bar{Q}$. This set is weakly increasing in $\bar{Q}$ and increases until the set of stochastically stable states is $\{\omega(q) : q \in [q^c, q^w]\}$, which happens at $q^w - q^c = \frac{\bar{Q}}{n}$. For sufficiently large $\bar{Q}$ we are thus back to the prediction with global imitation. On the other hand, as $\bar{Q}$ becomes very small the set of stochastically stable states approaches $\omega(q^w)$. The outcome thus continuously becomes more competitive as firms become less willing to imitate across markets.

Imitate the Best Average

It turns out that there are many similarities between how the dynamics work in the case of IBM and IBA. As when agents follow IBM, all absorbing states in which quantities not in $[q^c, q^w]$ are produced can be left with single experimentations. All polymorphic absorbing states can be left with chains of single experimentation transitions that bring previously separated markets within comparison reach. This would lead us to believe that $\omega(q^s)$ is a good candidate for stochastic stability. However,

[13] In Appendix B we give an informal argument of why this happens.

there is a problem: if there are single experimentations from $\omega(q^s)$ that are profitable in relative terms and bring the experimenting market out of comparison reach, then not even $\omega(q^s)$ is stable against single experimentations. This happens if $q^w - \frac{\bar{Q}}{n} > q^s$, which always holds for sufficiently small $\bar{Q}$. We can thus straightforwardly derive the result that $\omega(q^s)$ is stochastically stable as long as $\bar{Q}$ is sufficiently large, but it's more difficult to conclude what happens for small $\bar{Q}$. For example, the quantities $[q^s, \varphi(q^s)]$ could be conjectured to be stochastically stable, since they can all be connected by 1-experimentation transitions. However, from states $(q^s, \omega(\varphi(q^s)))$ there are states $\omega(q)$, $q < q^s$ that can also be reached by 1-experimentation transitions (see Figure 2.2). From these even larger quantities are profitable in relative terms. In the end, all states in Ω_M may turn out to be stochastically stable. In this case the concept of stochastic stability provides no prediction at all. We provide the following partial characterization of the stochastically stable states when IBA is used:

Proposition 2.10. *If $\bar{Q} \geq q^w \frac{n(k-1)}{(n+1)k-2}$, then $\omega(q^s)$ is the unique stochastically stable state.*

Proof. Step 1: All states in Ω_P can be left for states in Ω_M through 1-experimentation transitions.

This follows directly from step 1 in the proof of the preceding Proposition.

Step 2: $\omega(q^s)$ can be reached through a series of 1-experimentation transitions from any state in $\Omega_M \backslash \omega(q^s)$.

Take any state $\{\omega(q) \in \Omega_M \backslash \omega(q^s)\}$ and let some firm ij experiment in the direction of q^s to $\min\{q + \bar{Q}, q^s\}$ if $q < q^s$ and to $\max\{q - \bar{Q}, q^s\}$ if $q > q^s$. All firms imitate the experimentation with pp. Proceed in this way until $\omega(q^s)$ is reached.

Step 3: *$\omega(q^s)$ cannot be left with a single experimentation as long as $\bar{Q} \geq q^w \frac{n(k-1)}{(n+1)k-2}$.* The only way for an experimentation to q' from $\omega(q^s)$ to be imitated is if it brings the experimenting market out of comparison reach and at the same time $d(q^s, q') \geq 0$. Such an experimentation is possible if and only if

$q^w - \frac{\bar{Q}}{n} > q^s$ (Step 3 of the proof of the previous Proposition) which is equivalent to $q^s = q^w \frac{nk-1}{k(n+1)-2} \geq q^w - \bar{Q}/n \leftrightarrow \bar{Q} \geq q^w \frac{n(k-1)}{(n+1)k-2}$.

Step 4: Given step 3 we obtain $R(\omega(q^s)) > 1$ and given step 2 $CR^*(\omega(q^s)) = 1$, which means that $\omega(q^s)$ is stochastically stable. □

The quantity q^s thus conserves some of its stability properties also in the case of endogenous local imitation. However, it requires that firms aren't too cautious about imitating firms in other markets. When markets become too isolated, it seems that it looses its attraction as the unique stable quantity. The intuition is that when markets are more isolated, experimentations are more likely to make two previously similar market differ, thereby inducing them to evolve in isolation, which makes the calculations of average profits of strategies loose its force.

2.5.2 Markets Arranged Around a Circle

In this section we restrict the amount of information a firm has access to in an alternative way. We consider markets that are arranged around a circle and assume that each firm has access to information about quantities and profits of firms in neighboring markets. Information is therefore again local, but in contrast to the preceding section localness is exogenously imposed rather than assumed to depend on differences in the aggregate. An interpretation is that the geographic locations of the markets are such that not all firms can observe each other. An important implication of this information structure is that the same information will not be available in all the markets and that successful strategies will be bound to spread in a stepwise fashion over the population.

The Model

As previously, there are k markets, but these are now arranged around a circle and firms in each market only observe the quantities and profits of the firms in the 2η, $\eta < (k-1)/2$ neighboring markets. This creates a sense of localness of information. If

$\eta = \lceil (k-1)/2 \rceil$ [14], then firms have full information and we are back to the standard case. We refer to the set of firms in markets $(j-\eta), ..., (j-1), (j), (j+1), ..., (j+\eta)$ as the neighborhood of the firms in market j, and denote this set $\eta(j)$. Note that the markets in $\eta(j)$ also corresponds to the set of markets that observe j. The following represents a straightened segment of the circle of markets:

$$..., (j-\eta-1)\underbrace{(j-\eta), ..., (j-1), (j), (j+1), ..., (j+\eta)}_{\text{Markets observed by firms in market } j.}(j+\eta+1), ...$$

We will assume that there is some inertia, which means that we do not allow for the case in which all firms always adjust strategies in a perfectly synchronized way. We will show that with this assumption the main results of the paper are robust to this setting of exogenous local imitation. With no inertia, this is not the case. In this section we do not restrict ourselves to the linear case, so demand and cost functions are as in the standard setup.

Imitate the Best Max

First we note that as in the full information setting, the set of absorbing states coincides with Ω_M. We have:

Lemma 2.3. *When firms use IBM, there are k markets with n firms in each arranged around a circle and firms observe the $\eta < \lceil (k-1)/2 \rceil$ neighboring markets, the set of absorbing states is Ω_M.*

Proof. Evidently, all states in Ω_M are absorbing.

It remains to be shown that only states in Ω_M are absorbing. First, note that states in which $\exists ij, i'j$ such that $q_{ij} \neq q_{i'j}$ are left with some probability for states in which $q_{ij} = q_{i'j}$ for all $ij, i'j$. Next, suppose the highest profit in all markets is

[14] By $\lceil x \rceil$ we mean the smallest integer larger than x.

obtained in market j, where q^* is produced. Then all firms in $\eta(j)$ imitate q^* with pp. Next, all firms in $\eta(j+\eta)$ and $\eta(j-\eta)$ imitate q^* with pp. Iteration of this argument leads us to a state in which q^* is produced in all markets. $\square$

The characterization of the stochastically stable states coincides with that of the full information setting:

Proposition 2.11. *When firms use IBM, there are k markets with n firms in each arranged around a circle and firms observe the $\eta < \lceil (k-1)/2 \rceil$ neighboring markets, the set of stochastically stable states is $\{\omega(q) : q \in [q^c, q^w]\}$.*

Proof. Step 1: *A single experimentation in market j is imitated with pp. by the firms in markets $\eta(j)$ if and only if $D(q,q') := \pi(q',(n-1)q+q') - \max\{\pi(q,(n-1)q+q'), \pi(q,nq)\} \geq 0$.*

Step 2: *A double experimentation in market j is imitated with pp. by the firms in $\eta(j)$ if and only if $\pi(q',(n-2)q+q'+q'') - \max\{\pi(q,(n-2)q+q'+q''), \pi(q,nq)\} \geq 0$.*

Step 1 and step 2 imply that a strategy spreads with pp. to neighboring markets only if it spreads to the whole population with pp. in the full information case. In other words, this is a necessary condition for an experimentation to spread.

Step 3: *Single experimentations and double downward experimentations (as those in Proposition 2.1) spread to the whole population with pp. in the circular framework if they spread to the whole population in the global information framework*:

We consider single experimentations from $\omega(q)$ to q' and the upward double experimentations from $\omega(q)$ to q' *capable of causing a transition to $\omega(q')$ in the global information setup.*

(1) $\omega(q) \xrightarrow{1} \omega(q')$, $\omega(q) \xrightarrow{2} \omega(q')$, $q < q'$: (i) Let ij experiment to q'. (ii) Let one firm in all markets in $\eta(j)$ imitate. (iii) Let one firm in all markets in $\eta(j+\eta)$ and $\eta(j-\eta)$ imitate q' and proceed in this way until in all markets there are $n-1$ or $n-2$ firms producing q and 1 firm producing q'. (iv) Let all firms imitate q'.

(2) $\omega(q) \xrightarrow{1} \omega(q')$, $q' < q$: (i) Let ij experiment to q'. (ii) Let one firm in all markets in $\eta(j)$ imitate to q'. (iii) Proceed in this way until in all markets there are $n-1$ firms producing q and 1 firm producing q'. (iv) Let all firms imitate q'.

In other words, if a single or double experimentation spreads to the entire population in the standard setting it spreads to the entire population in the circular setting. Single and double experimentations then spread to the entire population in the circular setting if and only if they do so in the full information setting, and the set of stochastically stable states are the same in both cases (recall that in the proof only single and double experimentations are used to prove stochastic stability). □

The proof of Proposition 2.11 depends on the assumption of some inertia. In particular, we "freeze" the process while the invading strategy spreads to the entire population, though we do so in a more extreme way than necessary. For example, for upward single experimentations to quantities $q' \leq q^w$ it is sufficient that each time the strategy spreads, it does not spread to an entire market at once. It is worth to note that in the models considered in previous sections, inertia does not affect the set of stochastically stable states, whereas in this setting it becomes important. A reason for this is that when information is local, the way in which strategies spread is restricted by the imposed structure of localness. In some sense, in models of local imitation, there is some innate inertia, in the sense that what happens in one market can come to other markets' attention only in a restricted, structured way. Adding inertia to the model looses up some of this structure. For example, an experimentation can spread to all markets before an entire market switches, which is what is used in the proof of Proposition 2.11.

Imitate the Best Average

Also in this case, the result is closely related to the main result of the full information case. The set of absorbing states is again Ω_M. The proof is completely analogous to the case of IBM, so we do not repeat it here. We obtain:

Proposition 2.12. *The state $\omega(q^s(2\eta+1,n))$ is the unique stochastically stable state when there are k markets with n firms in each arranged around a circle, firms use IBA and observe the $\eta < \lceil (k-1)/2 \rceil$ neighboring markets.*

Proof. Step 1: *No single experimentation is imitated with pp. when the system is in $\omega(q^s(2\eta+1,n))$.*

This is simply since each firm observes $2\eta+1$ other markets and the quantity $q^s(2\eta+1,n)$ therefore is evolutionary stable with respect to average payoffs (see Proposition 2.3).

Step 2: $\omega(q) \xrightarrow{1} \omega(q^s(2\eta+1,n))$ *is possible from any state $\omega(q)$.*

We know that $\omega(q^s(2\eta+1,n))$ is a global invader, which means that any firm in $\eta(j)$ imitates an experimentation to $q^s(2\eta+1,n)$ with pp. starting in any state $\omega(q)$. We consider upward and downward experimentations to q^s:

1. $\omega(q) \xrightarrow{1} \omega(q^s(2\eta+1,n))$, $q < q^s(2\eta+1,n)$: (i) Let a firm ij experiment to $q^s(2\eta+1)$. (ii) Let one firm in each of the markets in $\eta(j)$ imitate $q^s(2\eta+1,n)$. Now, all firms in markets $\eta(j\pm\eta)$ still producing q will observe $x \in \{1,..,2\eta\}$ markets in which one firm produces $q^s(2\eta+1,n)$ and $n-1$ firms produce q, and $2\eta+1-x$ markets in which all firms produce q. They then imitate $q^s(2\eta+1,n)$ with pp. iff $\pi(q^s(2\eta+1,n),(n-1)q+q^s(2\eta+1,n)) - (\frac{(n-1)x}{n(2\eta+1)-x}\pi(q^s(2\eta+1,n),(n-1)q+q^s(2\eta+1,n)) + \frac{n(2\eta+1-x)}{n(2\eta+1)-x}\pi(q,nq)) \geq 0$. But this expression corresponds to $f(q,q')$ with a "lambda" equal to $\lambda' = \frac{(n-1)x}{n(2\eta+1)-x} > \lambda(2\eta+1,n)$. We know that q^s increases in λ, which means that (abusing notation) $q^s(\lambda') > q^s(2\eta+1,n)$. Then, by Corollary 2.2 firms in $\eta(j\pm\eta)$ imitate $q^s(2\eta+1,n)$ with pp. So let one firm in the markets in $\eta(j\pm\eta)$ producing q imitate $q^s(2\eta+1,n)$. Proceed in this way until in all markets there are $n-1$ firms producing q and 1 firm producing $q^s(2\eta+1,n)$. (iii) Let all the remaining firms imitate $q^s(2\eta+1,n)$, which they do with pp. since $q^s(2\eta+1,n)$ is in the direction of q^w and therefore has a relative payoff advantage in all markets.

2. $\omega(q) \xrightarrow{1} \omega(q^s(2\eta+1,n))$, $q > q^s(2\eta+1,n)$: (i) Let ij experiment to $q^s(2\eta+1,n)$. (ii) Let the remaining firms in j and all firms in $\eta(j)$ in imitate to $q^s(2\eta+1,n)$. (iii) Since all firms in the experimenting markets now produce lower quantities they obtain higher profits, since $q^s(2\eta+1,n)$ is greater than the "monopoly" outcome,

and are thus imitated with pp. by all firms in $\eta(j \pm \eta)$. Let these firms imitate $q^s(2\eta + 1, n)$. (iv) Proceed in this way until all firms produce $q^s(2\eta + 1, n)$.

Since $\omega(q) \xrightarrow{1} \omega(q^s(2\eta+1, n))$ and $\omega(q^s(2\eta+1, n)) \xrightarrow{\geq 1} \omega(q)$, for all $q \in \Gamma$, we have that $R(\omega(q^s(2\eta+1))) > 1$, $Cr^*(\omega(q^s(2\eta+1))) = 1$ and therefore $\omega(q^s(2\eta+1, n))$ is stochastically stable. $\square$

Proposition 12 shows that q^s conserves its stability properties when imitation is not completely global and strategies can spread only in a stepwise process across the population. It is worth noting that the outcome becomes more competitive in this setting, in which information is more restricted compared to the case of global imitation. The result also becomes more competitive as η decreases and the set of information available to firms becomes smaller (by the comparative statics results on q^s). Less information about what happens in other market thus means more competition.

2.6 Concluding Remarks

It has been shown that outcomes tend to become less competitive when firms have a greater tendency to imitate across markets. In the benchmark setting, when firms use IBM the whole interval between the Cournot and Walrasian outcome becomes stochastically stable. This result corresponds to the outcome of the single market memory model of Alós-Ferrer (2004), which points to the close relationship between such models and imitation in settings with multiple markets. When IBA is used a distinct feature of the outcome is that it depends on the number of markets. This means that if firms behave according to IBA, counter-intuitive things could happen. An additional market may actually decrease competition. This would be the case if for example there is a single local market for a certain product in a city and firms open up and start providing the same product in another city. If firms imitate across markets this may then actually decrease competition.

The uniqueness of q^s and the fact that it is related to the number of markets, a parameter that is of no relevance in standard Cournot oligopolies, makes it suitable for experimental testing. If variations of the number of markets in an oligopoly experiment in which individuals observe outside markets have no effect on the outcome, then this would be evidence against IBA and/or imitation across markets. This paper thus provides a new prediction that can be used to test if individuals imitate across markets and if they do so according to IBA.

We have studied two alternative informational settings in which imitation is local in different ways. It has been shown that the cautiousness with respect to imitating firms in markets with different aggregates affects the long run outcome. More cautious firms leads to more competitive results. The firms thus benefit in the long run by not being too sensitive about imitating firms in markets with different aggregates. Apesteguia et. al. (2007) obtain a result of similar flavor. In their model, imitation of only the non-competitors leads to Nash equilibrium, whereas imitation of only the competitors leads to the Walrasian outcome. The model of endogenous local imitation allows firms to *sometimes* imitate across markets. In this setting, the outcome continuously becomes more competitive as firms become more willing to imitate non-competitors. A similar conclusion is obtained when firms are located around a circle. In this case, more information about other markets leads to less competitive results if firms imitate according to IBA.

A conclusion that holds across these models is that more information about firms in other markets tends to lead to less competitive outcomes. This should be contrasted with the conclusion of Huck et. al. (1999, 2000). These authors conclude that more information *about the firms in the own market* leads to more competitive results and relate this to the prediction of Vega-Redondo (1997). Whereas they conclude that their results can be taken as tentative evidence in favor of publishing information about firm performance, the results of the present paper gives the contrary indication. If firms imitate non-competitors, publishing this information could actually lead to less competition.

In this paper markets have been assumed to be ex-ante identical. Having a set of ex-ante identical markets is a strong assumption, but it helps in obtaining sharp analytical results. At the same time imitation is more intuitive when there are no structural differences between the environment of the imitator and the sampled firms. However, an interesting direction for further research would be to consider markets that are ex-ante different in some dimension.

There is a close relationship between multi-market models and models with memory. In the memory models of Alós-Ferrer (2004) and Bergin and Berghardt (2009) an assumption of no inertia is needed for the obtained results. In the multimarket model, inertia becomes important when we consider markets arranged around a circle. An interesting avenue for further research is to analyze more closely the relationship between memory, local interaction and inertia.

Appendix A: Methods for Proving Stochastic Stability

We use three different methods in order to prove stochastic stability: (i) the "tree-surgery" method of Young (1993, 1998), (ii) the radius-coradius argument by Ellison (2000) and (iii) the modified radius-coradius of Ellison (2003). We here briefly summarize these methods when applied to a process such as the one outlined in Section 2.

(i) Young's (1993, 1998) tree surgery argument: For any $\omega \in \Omega$, an ω-tree is a tree branching out from ω and reaching every other absorbing state ω' via a unique path directed from ω' to ω. The cost $c(\omega', \omega)$ of an edge between nodes ω' and ω in this graph is equal to the minimum number of experimentations required for the imitation dynamics to move the system from ω' to ω with positive probability. Let Θ_ω be the set of all ω-trees and let $\theta^*_\omega = \underset{\theta \in \Theta_\omega}{\arg\min} \sum_{(\omega', \omega'') \in \theta} c(\omega', \omega'')$. Then the stochastic potential of state ω is defined as $\gamma_\omega = \sum_{(\omega', \omega'') \in \theta^*_\omega} c(\omega', \omega'')$. In other words, the stochastic potential of ω is equal to the sum of the costs of all the edges of the minimum cost ω-tree. The set of stochastically stable states is $\underset{\omega \in \Omega}{\arg\min}\{\gamma_\omega\}$, i.e. the states with the smallest stochastic potential.

(ii) Ellison's (2000) radius-coradius argument: For any $\Omega_r \subset \Omega$, let the radius of Ω_r, denoted by $R(\Omega_r)$, be the minimum number of experimentations needed for the imitation dynamics to lead the system out of Ω_r to some state in $\Omega \backslash \Omega_r$ with positive probability. Let the coradius of Ω_r, denoted by $Cr(\Omega_r)$, be the maximum (over all $\omega' \in \Omega \backslash \Omega_r$) number of experimentations needed to enter Ω_r from $\Omega \backslash \Omega_r$. Then a sufficient condition for the set of stochastically states to be contained in Ω_r is that $R(\Omega_r) > Cr(\Omega_r)$. The expected waiting time to reach Ω_r if the process starts in $\Omega \backslash \Omega_r$ is in the order of magnitude of $\epsilon^{-Cr(\omega)}$.

(iii) Ellison's (2000) modified radius-coradius: Let $(\omega^1, \omega^2, ..., \omega^T)$ represent a path originating in some state $\omega^1 \in \Omega \backslash \Omega_r$ and ending up in some state $\omega^T \in \Omega_r \subset \Omega$ and where $\omega^i \in \Omega \backslash \Omega_r$ $\forall i = 2, ..., T-1$. Let $c(\omega^i, \omega^{i+1})$ be the minimum number of experimentations needed for the imitation process to bring the system from ω^i to ω^{i+1} with positive probability. Let $c(\omega^1, \omega^2, ..., \omega^T) = \sum_{i=1}^{T} c(\omega^i, \omega^{i+1}) - \sum_{i=2}^{T-1} R(\omega^i)$. Let $S(\omega^1, \Omega_r)$ be the set of paths originating in ω^1 and ending up in some state in Ω_r. Let $c^*(\omega^1, \omega^2, ..., \Omega_r) := \min\limits_{(\omega^1, \omega^2, ..., \omega^T) \in S(\omega^1, \Omega_r)} c(\omega^1, \omega^2, ..., \omega^T)$. The modified coradius of Ω_r is then $Cr^*(\Omega_r) := \max\limits_{\omega^1 \in \Omega \backslash \Omega r} \{c^*(\omega^1, \omega^2, ..., \Omega_r)\}$. If $R(\Omega_r) > Cr^*(\Omega_r)$, then the set of stochastically stable states is contained in Ω_r.

Appendix B: Non-Stability of Polymorphic States

Consider the model analyzed in Section 5.1. Assume that there is some capacity constraint $\delta v > q^w$. The argument that we used to destabilize polymorphic states with single experimentations depended on experimentations to large enough quantities to bring two separated markets into comparison reach. Here we give an argument that shows that in the linear setting even with a capacity constraint polymorphic states tend to be unstable.

We know that any state (and this holds for both polymorphic and monomorphic states) in which quantities in $\Gamma \backslash [q^c, q^w]$ are produced are inherently unstable. For polymorphic states to have any hope of being stochastically stable they should therefore involve only quantities in $[q^c, q^w]$. Now consider a polymorphic state in which $q^c = \frac{a-c}{b(n+1)}$ and $q^w = \frac{a-c}{bn}$ are produced in markets j and j', implying $Q_j = \frac{n(a-c)}{b(n+1)}$ and $Q_{j'} = \frac{a-c}{b}$. Then, an experimentation of one firm in j' to $q' = 0$ means we will now have $Q_{j'} = \frac{(a-c)(n-1)}{bn}$ and this is lower than Q_j. Hence there are downward experimentation available for firms in j' such that j and j' will be within comparison reach and they will therefore imitate the same quantity with positive probability. The same argument can be made for any two initial quantities in $[q^c, q^w]$. Any polymorphic state in which quantities in $[q^c, q^w]$ are produced can therefore be destabilized with a single experimentation. By iterating such experimentations a monomorphic state can be reached and arguments similar to those in the proof of Proposition 2.9 can be used to show that no polymorphic state is stochastically stable.

Appendix C: Models with Memory as Multimarket Models

Single market models with memory, as in Alós-Ferrer (2004) and Bergin and Bernhardt (2009) can be seen as special cases of multi-market models. The additional markets can be thought of as existing in the memory of agents. In the models of Alós-Ferrer (2004) and Bergin and Bernhardt (2009), it is assumed that agents remember quantities and profits k periods back, creating a total of $k+1$ markets. They then imitate according to either IBM or IBA observing quantities and profits in the present and in the past. There is no exogenous inertia in the models, so all firms change strategies in perfect synchronicity. As in the multi-market model $\{\omega(q) : q \in \Gamma\}$ constitute the set of absorbing states and the set of stochastically stable states is contained in this. However, whereas in a multi-market model strategies can spread to new markets in an arbitrary order, in memory models strategies cannot spread backwards in time. This creates a special kind of inertia. If we think of the memory model as $k+1$ markets, where market 1 is the present, market 2 is what happened yesterday and so on, another observation is that experimentations can occur only in market 1. If we consider some state $\omega(q)$, a single experimentation in market 1 in period t will be imitated in a very specific sequence. This sequence of imitation can be thought of as proceeding from the most distant past, market $k+1$, successively to the present. In $t+1$ all firms in market $k+1$ are given the possibility to imitate, in $t+2$ all firms in market k are given the possibility to imitate and so on until the turn comes to market 1. The different markets are thus picked to imitate one at a time in a determined sequence.

The crucial question for the relationship between the results of the multi-market model with the memory model is how this inertia affects the results. A first answer to this question is that any destabilization achieved in a memory model is valid also in a multi-market model with inertia. This is simply because the presence of inertia enables us to clone any sequence of imitations in the memory model. However, with arbitrary inertia in the multi-market model, imitation can occur in sequences not possible in the memory model. This means that there are more ways in which a state can be destabilized. As it turns out, this does not affect the result when IBM is used (as seen in Proposition 2.1), but the results are completely different when IBA is used. In the latter case, the collusive outcome becomes increasingly difficult to destabilize as memory becomes longer, which eventually makes it stochastically stable (Bergin and Bernhardt (2009)).

Bibliography

[1] Apesteguia, J., Huck, S. and Oechssler, J. (2007). Imitation: Theory and Experimental Evidence. *Journal of Economic Theory 135*, 217-235.

[2] Apesteguia, J., Huck, S., Oechssler, J. and Weidenholzer, S. (2010). Imitation and the Evolution of Walrasian Behavior: Theoretically Fragile but Behaviorally Robust. *Journal of Economic Theory 145*, 1603-1617.

[3] Apesteguia, J. and Selten, R. (2005). Experimentally Observed Imitation and Cooperation in Price Competition on the Circle. *Games and Economic Behavior 51*, 171-192.

[4] Alós-Ferrer, C. (2004). Cournot vs Walras in Dynamic Oligopolies with Memory. *International Journal of Industrial Organization 22:2*, 193-217.

[5] Alós-Ferrer, C. and Ania, A.B. (2005). The Evolutionary Stability of Perfectly Competitive Behavior. *Economic Theory 26*, 497-516.

[6] Bergin, J. and Bernhardt, D. (2009). Cooperation through Imitation. *Games and Economic Behavior 67*, 376-388.

[7] Bosch-Domenech, A. and Vriend, N.J. (2003). Imitation of successful behavior in Cournot markets. *Economic Journal 113*, 495-524.

[8] Conlisk, J. (1996). Why Bounded Rationality?. *Journal of Economic Literature 34*, 669-700.

[9] Ellison, G. (2000). Basins of Attraction, Long-Run Stability, and the Spread of Step by Step Evolution. *The Review of Economic Studies 67*, 17-45.

[10] Eshel I., Samuelson L. and Shaked A. (1998). Altruists, Egoists and Hooligans in a Local Interaction Model. *American Economic Review 88*, 157-179.

[11] Freidlin, M. and Wentzell, A.D. (1984). Random Perturbations of Dynamic Systems. New-York: Springer-Verlag.

[12] Fudenberg, D. and Levine, D. (1998). The theory of Learning in Games. MIT press, Cambridge MA.

[13] Gigerenzer, G., Todd, P.M. and the ABC Research Group. (1999). Simple Heuristics that Make us Smart. New York. New York: Oxford University Press.

[14] Huck, S., Normann, H.T, and Oechssler, J. (1999). Learning in Cournot Oligopoly: an Experiment. *Economic Journal 109*, 80-95.

[15] Huck, S., Normann, H.T, and Oechssler, J. (2000). Does Information about Competitors Actions Increase or Decrease Competition in Experimental Oligopoly Markets?. *International Journal of Industrial Organization 18*, 39-57.

[16] Jun T. and Sethi R. (2007). Neighborhood Structure and the Evolution of Cooperation. *Journal of Evolutionary Economics 17*, 623-646.

[17] Kandori, M., Mailath, G. and Rob, R. (1993). Learning, Mutation and Long Run Equilibria in Games. *Econometrica 61*, 29-56.

[18] Mengel, F. (2009). Conformism and Cooperation in a Local Interaction Model. *Journal of Evolutionary Economics 19*, 397-415.

[19] Offerman, T., Potters, J. and Sonnemans, J. (2002). Imitation and Belief Learning in an Oligopoly Experiment. *Review of Economic Studies 69*, 973-997.

[20] Offerman, T. and Schotter, A. (2009). Imitation and Luck: An Experimental Study of Social Sampling. *Games and Economic Behavior 65*, 461-502.

[21] Pingle, M. and Day, R.(1996). Modes of Economizing Behavior: Experimental Evidence. *Journal of Economic Behavior and Organization 29*, 191-209.

[22] Schenk-Hoppé, KR. (2000). The Evolution of Walrasian Behavior in Oligopolies. *Journal of Mathematical Economics 33*, 35-55.

[23] Schaffer, M. (1988). Evolutionary Stable Strategies for a Finite Population and a Variable Contest Size. *Journal of Theoretical Biology 32*, 469-478.

[24] Schipper, B. (2008). Imitators and Optimizers in a Cournot Oligopoly. University of California Davis, Department of Economics, Working Paper.

[25] Shapiro, C. (1989). Theories of Oligopoly Behavior. "Handbook of Industrial Organization". ed. Schmalensee, R. and Willig, R.D. North Holland. Vol. 1, Chapter 6, pp. 330-414.

[26] Tanaka, Y. (1999). Long run equilibria in an asymmetric oligopoly. *Economic Theory 14*, 705-715.

[27] Tanaka, Y. (2000). Stochastically Stable States in an Oligopoly with Differentiated Goods: Equivalence of Price and Quantity Strategies. *Journal of Mathematical Economics 34*, pp. 235-253.

[28] Vega-Redondo, F. (1997). The Evolution of Walrasian Behavior. *Econometrica 65*, 375-384.

[29] Vega-Redondo, F. 2003. Economics and the Theory of Games. Cambridge University Press, Cambridge.

[30] Young, P. 1993. The Evolution of Conventions. *Econometrica 61*, 57-84.

[31] Young, P. 1998. Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. New-Jersey Princeton University Press.

.

Chapter 3

Altruistic Provision of Public Goods and Local Interaction

3.1 Introduction

A standard assumption in economic theory is that individuals act in their own best interests. Nevertheless, people are often observed acting beyond their immediate self-interests, which is commonly referred to as *altruism*. There are several approaches to reconcile such acts with the rationality typical of many economic models. One is to incorporate the well being of others into individuals' utility functions (e.g. Becker, 1974 and 1981). Another is to consider infinitely repeated interaction, in which case trigger strategies that punish selfish behavior can sustain altruism as a Nash equilibrium (see, for example, Fudenberg and Maskin 1986). Whereas these (by now standard) approaches show that altruism can be consistent with rationality, altruistic behavior can also arise if choices are made in a boundedly rational way. Eshel, Samuelson and Shaked (1998, henceforth, ESS) consider a model in which individuals live on a circle and repeatedly choose whether to provide a local public good. Instead of the standard assumption of rationality, individuals make their choices by imitating successful neighbors. In this setting altruism can persist and even coexist with selfish behavior. The intuition is that the local interaction structure makes it possible to form local altruistic communities where payoffs are high, which leads altruistic choices to be imitated.

Both the assumption of local imitation and of a local externality[1] are reasonable in many settings. For example, several experiments document the importance of imitative behavior in economic situations (e.g. Huck, Normann and Oechssler 1999, Selten and Apesteguia 2005, Apesteguia Huck and Oechssler 2007).[2] There are also theories in psychology which argue that the bulk of human behavior is learned by observation and imitation of others (see for example Bandura (1977)). Imitation is likely to be local, in particular if it is *payoff biased*, in the sense that strategies generating higher payoffs are imitated more frequently. In this case, an individual must observe both the actions and consequential well being of his potential role models. This is quite a demanding informational requirement that probably is met mainly when it comes to family, friends and closer acquaintances. Moreover, in the theory of social groups in the sociology literature, it is argued that individuals are most strongly influenced by members of their primary groups, which consist of people they interact with frequently, such as family, friends, colleagues and neighbors[3]. It is also evident that many externalities are local. For example, overconsumption of subsidized water affects the local water supply and possibly the local reserves of groundwater, creating a negative local externality. Fisheries along the coast affect the supply of fish mostly in a local or regional way. Many pollutants, such as particles, lead, sulphur and nitrogen oxides act in a local or regional way and can be related to such everyday decisions as deciding what kind of lawn mower to buy (where the noise level is also an issue). The externalities caused by the decision to drive a car and how to do it, such as traffic, accidents and pollution (except CO2) are primarily non-global. Littering close to where we live affects mainly people living in the same area, as does letting the dog run loose.[4]

However, these examples also point to a drawback of the framework of ESS. In their model, the externality affects only the two closest neighbors, whereas in

[1] Providing a public good can equivalently be thought of as choosing an action with a positive externality, or not choosing one with a negative externality.

[2] Some papers do not find imitation significant, e.g. Kirchkamp and Nagel (2007).

[3] See for instance Light, Keller and Calhoun, Sociology, 5th edition 1991, ch. 8.

[4] For more information on some of the externalities mentioned in this paragraph, see Tietenberg & Lewis, (2009).

the mentioned examples it is likely to affect a much larger number of individuals. This problem has been considered in the literature. Jun and Sethi (2007), Matros (2008), and Mengel (2009) show that altruism can persist when the externality affects exactly the same (possibly large) set of individuals that each individual imitates. However, in many situations it will not hold that the set of individuals that are affected by the externality is identical to the set of potential role models. This issue is partially considered by Mengel (2009), who shows that altruism will not survive when imitation is less local than the externality. However, it is likely that in many cases precisely the opposite is true. While the external effects of our actions often (as in the examples mentioned above) affect a large number of individuals, we mostly learn behaviors from a more limited set of individuals. For instance, we probably learn to be conservative in water consumption by observing the behavior of those close to us, whereas our actions affect a much larger set of individuals. The existing literature, however, provides no result for this case. This means that we do not know to what extent local externalities combined with local imitation can explain altruistic behavior when the externality is less local than imitation. It is therefore important to see what results follow when taking this natural pairing of assumptions into account.

In this paper, I therefore extend the framework of ESS to study the case in which the externality is less local than imitation. As in ESS, individuals live on a circle and repeatedly decide whether to provide a public good[5] (to be an altruist), or not do it (and be an egoist). The public good is shared by an arbitrary number of neighbors. To make a decision, the individual observes the actions and payoffs of the two closest neighbors and imitates the action that generated the highest average payoff. The individuals also sometimes experiment, in which case they choose an action randomly. It is shown that in the absence of experimentation, altruism can persist and coexist with selfish behavior as long as there are at least two individuals in the population with which the public good is not shared. I.e. altruism can survive as long as the public good is non-global. With experimentation there is

[5] I.e. carry out an action with a positive externality.

an important interplay between the localness of the externality and the size of the population. Altruism can survive even if the externality reaches a larger number of individuals than those that are imitated, but only if the population is sufficiently large. The required size of the population increases in the number of individuals affected by the externality. Hence, altruism can persist both in the presence and in the absence of experimentation. The conclusion therefore is that local interaction and imitation is a possible explanation for altruistic behavior in situations well beyond those considered in ESS.

The intuition behind the results is that since the public good is local, altruists can group together and exclude egoists from their contributions. In this way, altruists have mainly altruist neighbors, and egoists have mainly egoist neighbors. Therefore, the altruists obtain higher payoffs and thus tend to be imitated. As the public good becomes less local, altruists need to form larger groups in order to exclude the egoists. When the public good is nearly global, at most one such altruist group can fit in the population. It turns out that such a constellation is very sensitive to experimentation and it is actually enough that a single altruist switches for the entire population to descend into egoism. A large population protects against this eventuality by allowing either several altruist groups or few but very large groups. These constellations are more robust to experimentation, since pockets of altruists can survive egoistic experimentations. For this reason, large populations help altruism to persist when the public good is less local.

The outline is the following: In section 2 the model is presented. Section 3 and 4 contain the characterization of the main results. Section 5 considers an extension and Section 6 concludes.

3.2 The Model

Consider a set of $N := \{1, 2, \ldots, n\}$ individuals that live on fixed locations around a circle. The immediate neighbors to the left and right of $i \in N$ are denoted $i-1$ and $i+1$ respectively. The second neighbors to the left and right of i are denoted $i-2$ and $i+2$, and so on. The $2z$ closest neighbors of i thus consist of $\{i-z, i-z+1, ..., i-1, i+1, ..., i+z-1, i+1\}$.

The model proceeds in discrete time. At each time period $t \in \{1, 2, ...\}$ each $i \in N$ is drawn with independent and identical probability $\mu \in (0, 1)$ to choose an action from $S := \{a, E\}$.[6] Once a choice is made, it remains the same until the individual again is drawn to revise his choice. Let $s_{it} \in S$ be the choice of $i \in N$ in t. If $s_{it} = a$, then i provides a local public good in t. In this case, i incurs a cost $c < 1/2$ in t.[7] If $s_{it} = E$, no cost is incurred. Let N_i^E consist of i's $2n_E \geq 2$ closest neighbors (not including i). N_i^E is referred to as i's *externality neighborhood*. If $s_{it} = a$, then i contributes with 1 unit of utility to each $j \in N_i^E$ in t. If $s_{it} = E$, then i contributes with no utility to his neighbors in t. Hence, if $s_{it} = a$, i is said to be an *altruist* in t and if $s_{it} = E$, i is said to be an *egoist* in t. Note that being an altruist is a strictly dominated action, since the individual incurs a net private cost c when he is an altruist. However, choosing to be an altruist is good for society, since the contributions to other individuals sum at least 2, and $2 - c > 0$. Note also that if $s_{it} = a$, then i incurs the cost c and provides utility to his neighbors in each period until he is drawn for another revision opportunity, where a new choice is made. Let $(a, E) := (1, 0)$. The payoff of i in t is then $\pi_{it} := \sum_{j \in N_i^E} s_{jt} - s_{it}c$. Notice that since the payoff of an individual depends on the choices of his neighbors, individuals may obtain different payoffs even if they are choosing the same action.

Let N_i^I consist of i's $2n_I \geq 2$ closest neighbors and i himself. N_i^I is referred to as i's *imitation neighborhood*. When given a revision opportunity in t, i observes his own payoff and action in $t-1$ and the payoffs and actions of the individuals in N_i^I in $t-1$. He computes the average payoff of each observed action in N_i^I and chooses the action that generated the highest average payoff. If $s_{it} = s_{jt}$ for all $j \in N_i^I$, then i keeps s_{it}. The reason that this imitation rule is considered is that the same

[6] This means that there is some *inertia* in the revision of choices, in the sense that individuals do not necessarily revise their choices in every time period. This can reflect an unwillingness of the individuals to revise their choices too often. It also allows for the possibility that not all individuals revise their choices in perfect synchronicity.

[7] As in ESS, egoism will trivially result if $c > 1/2$. This will be shown in section 3. $c = 1/2$ is avoided since it leads to situations with draws.

action may generate different payoffs in the same time period. For example, $i-1$ may obtain a high payoff when choosing a, while $i+1$ obtains a low payoff from the same action. In this situation, it is not clear that it would be attractive to imitate a. By computing the average payoff of a, i takes into account both the low and high payoff of a.

In other words, the dynamics work as following. Say i is given a revision opportunity in t. He then observes the actions and payoffs of all individuals in N_i^I in $t-1$ and makes a choice accordingly. Suppose this leads him to choose $s_{it}=a$ and that the next revision opportunity arrives at $t+3$. In this case i incurs a cost c and contributes with one unit of utility to the individuals in N_i^E in t, $t+1$ and $t+2$. In $t+3$ he observes the actions and payoffs of the individuals in N_i^I in $t+2$ and chooses s_{it+3}. The following illustrates an individual i and his imitation and externality neighborhoods when $n_E=3$ and $n_I=1$:

$$..., i-4, \overbrace{i-3, i-2, \underbrace{i-1, i, i+1}_{N_i^I}, i+2, i+3}^{N_i^E \cup \{i\}}, i+4, ...$$

What determines the choice of i in t is whether the average payoff of the altruists in N_i^I is larger than the average payoff of the egoists in N_i^I. Let $d\pi_{it}$ denote the *difference* between the average payoffs of altruists and egoists in N_i^I in t. This means that if $d\pi_{it}$ positive, then i chooses to be an altruist in $t+1$, if given a revision opportunity. We can write

$$d\pi_{it} := \frac{1}{\sum_{j\in N_i^I} s_{jt}} \sum_{j\in N_i^I} s_{jt}\pi_{jt} - \frac{1}{n_I - \sum_{j\in N_i^I} s_{jt}} \sum_{j\in N_i^I} (1-s_{jt})\pi_{jt}.$$

The first term corresponds to the average payoff of the altruists in N_i^I in t and the second term gives the average payoff of the egoists in N_i^I in t. The imitation rule just described implies that if i is given a revision opportunity in t, then

$$s_{it} = \begin{cases} a \text{ if } d\pi_{it-1} > 0 \\ E \text{ if } d\pi_{it-1} < 0 \end{cases}$$ [8]

If $n_E = n_I = 1$ the model reduces to the one analyzed by ESS. If $n_E \leq n_I$ the model is similar to the one analyzed by Mengel (2009). Here, the focus is on $n_E > n_I$, i.e. the public good is less local than imitation. For tractability the main focus (Sections 3-4) is on $n_E > n_I = 1$, i.e. individuals imitate only the two closest neighbors. However, in Section 5 a partial result for $1 < n_E < n_I$ is provided. We also impose $n_E \leq (n-1)/2$, which simply means that the externality neighborhood is not larger than the entire circle of individuals.

The model described so far defines a finite Markov chain, in which a state is a specification of whether each individual in N is an altruist or an egoist.[9] The state space of this process is therefore $\{a, E\}^n$ and the transition probabilities depend on the imitation rule and μ. This Markov process is denoted Γ and referred to as the *unperturbed process*. Let $P^\tau(\omega, \omega')$ be the probability of reaching state $\omega' \in \{a, E\}^n$ from a state $\omega \in \{a, E\}^n$ in τ periods. An *absorbing state* of Γ is defined as a state $\omega \in \{a, E\}^n$ such that $P^1(\omega, \omega) = 1$. Hence, an absorbing state is a state that once entered cannot be left. Let Ω denote the set of absorbing states. The absorbing states of Γ are characterized in Section 3.

As in ESS individuals are also allowed to sometimes deviate from the imitation rule and experiment. In each t each individual experiments with independent and identical probability ε. If i experiments, he reverses his choice and picks the opposite of that prescribed by the imitation rule. By incorporating experimentation a different Markov process is obtained, which is referred to as the *perturbed process* and denoted Γ^ε. The experimentations make a transition between any two states of Γ^ε possible and the Γ^ε is therefore irreducible and aperiodic. By a well known results, this implies that Γ^ε has a unique stationary distribution which describes average behavior in the long run.[10] Denote this stationary distribution u_ε. Let $u^* := \lim_{\varepsilon \to 0} u_\varepsilon$,

[8] Ties.do not occur as long as $c \neq \frac{1}{2}$. This is implied by the proof of Lemma 3.1.
[9] A review on Markov processes can be found in Karlin and Taylor, 1975.
[10] See, e.g., Karlin and Taylor 1975.

i.e. u^* is the stationary distribution of Γ^ε as the experimentation probability approaches zero. The support of u^* is referred to as the set of *stochastically stable states*. Denote the set of stochastically stable states Ω^*. With vanishing experimentation probability the fraction of time spent in the stochastically stable states approaches one as time approaches infinity. By a by now standard result $\Omega^* \subset \Omega$.[11] As mentioned by Binmore, Samuelson and Vaughan (1995), the absorbing states of Γ give an idea of where the dynamics of Γ^ε end up in the short run, whereas the stochastically stable states of Γ^ε is where the process will spend most of its time in the very long run. Ω^* is characterized in section 4.

3.3 Absorbing States of the Unperturbed Process

I first show how the unperturbed process behaves. The first result, Lemma 3.1, will be useful in subsequent proofs and provides intuition concerning how the imitation dynamics work. In Lemma 3.1 and in what follows, time indexes will be omitted (whenever this does not cause confusion) in order to reduce notation. Let $\overline{s_i} := (s_{i-n_E-1}, s_{i-n_E}, s_{i+n_E}, s_{i+n_E+1})$. In other words, $\overline{s_i}$ specifies the actions chosen by the individuals just inside and just outside i's externality neighborhood.

Lemma 3.1. *Suppose $n_E > n_I = 1$. (i) $d\pi_i$ depends only on $\overline{s_i}, s_{i-1}, s_i, s_{i+1}$ and c. (ii) $d\pi i > 0$ if and only if either $(s_{i-1}, s_i, s_{i+1}) = (a, a, a)$ or $\overline{s_i} = (a, a, E, E)$, $c < 1/2$ and $(s_{i-1}, s_{i+1}) = (a, E)$ (or the mirror image of this).*

Proof. (i) We can write $d\pi_i = \delta(1/2(\pi_j + \pi_k) - \pi_l)$, where δ is an indicator function, taking the value 1 if there are two altruists in N_i^I and -1 if there is only one. j and k are the individuals in N_i^I picking the same action and obviously $i \in \{j, k, l\}$. Let $\eta_i := \sum_{j' \in N_i^E} s_{j'}$. Hence $d\pi_i = \delta(1/2(\eta_j + \eta_k) - \eta_l) - c$. Now, note that j, k, and l will all receive the contributions of the individuals in $\{i - n_E + 1, i - n_E + 2, \ldots, i + n_E - 2, i + n_E - 1\} \backslash \{j, k, l\}$. This means that these contributions will be added and subtracted once in $d\pi_i$ and they are therefore irrelevant for $d\pi_i$. Hence, $d\pi_i$ is a

[11] See, for example, Young 1993.

function only of the actions of individuals $\{j,k,l\}\cup\{i-n_{I-1},i-n_I,i+n_I,i+n_{I+1}\}$ and c.

(ii) If $(s_{i-1},s_{i,}s_{i+1})=(a,a,a)$ then $d\pi_i>0$ obviously. Likewise if $(s_{i-1},s_{i,}s_{i+1})=(E,E,E)$ then $d\pi_i<0$. Otherwise individual i is either surrounded by two individuals opposite of his kind, or by one of each kind. We therefore have the following four possibilities: Ea_iE, aE_ia, aa_iE, aE_iE. By eliminating irrelevant contributions and considering these four possibilities we obtain:

Ea$_i$**E**: $d\pi_i=\delta 1/2(\eta_j+\eta_k)-\delta\eta_l-c=\delta 1/2(s_{i-n_E-1}+s_{i-n_E}+s_i+s_{i+n_E}+s_{i+n_E+1}+s_i)-\delta(s_{i-n_E}+s_{i+n_E})-c=-1/2(s_{i-n_E-1}-s_{i-n_E}-s_{i+n_E}+s_{i+n_E+1})-c-s_i=-1/2(s_{i-n_E-1}+s_{i+n_E+1}-s_{i-n_E}-s_{i+n_E})-(c+1)$. The maximum of this expression is $-c-1$ and an isolated altruist will therefore always become an egoist, i.e. $d\pi_i<0$.

aE$_i$**a**: $d\pi_i=\delta 1/2(\eta_j+\eta_k)-\delta\eta_l-c=1/2(s_{i-n_E-1}+s_{i-n_E}+s_{i+1}+s_{i+n_E}+s_{i+n_E+1}+s_{i-1})-(s_{i-nI}+s_{i+nI}+s_{i-1}+s_{i+1})-c=1/2(s_{i-n_E-1}-s_{i-n_E}-s_{i+n_E}+s_{i+n_E+1})-c-\delta s_i=1/2(s_{i-n_E-1}+s_{i+n_E+1}-s_{i-n_E}-s_{i+n_E})-(c+1)$. The maximum of this expression is $-c$ and an isolated egoist will therefore always remain so, i.e. $d\pi_i<0$.

aa$_i$**E**: $d\pi_i=\delta 1/2(\eta_j+\eta_k)-\delta\eta_l-c=\delta 1/2(s_{i-n_E-1}+s_{i-n_E}+s_i+s_{i-n_E}+s_{i+n_E}+s_{i-1})-\delta(s_{i+n_E}+s_{i+n_E+1}+s_i+s_{i-1})-c=1/2(s_{i-n_E-1}+2s_{i-n_E}-s_{i+n_E}-2s_{i+n_E+1})-(c-s_i)=1/2(s_{i-n_E-1}+2s_{i-n_E}-s_{i+n_E}-2s_{i+n_E+1})-(c-1)$. This is positive only at the maximum, which is equal to $1/2-c$, and only if $c<1/2$. The maximum is attained at $\overline{s_i}=(a,a,E,E)$, hence $d\pi_i>0$ when $\overline{s_i}=(a,a,E,E)$, $c<1/2$, $(s_{i-1},s_{i+1})=(a,E)$ and $s_i=a$.

aE$_i$**E**:$d\pi_i=\delta 1/2(\eta_j+\eta_k)-\delta\eta_l-c=\delta 1/2(s_{i-n_E}+s_{i+n_E}+s_{i-1}+s_{i+n_E}+s_{i+n_E+1}+s_{i-1})-\delta(s_{i-n_E-1}+s_{i-n_E})-c=-1/2(-s_{i-n_E}-2s_{i-n_E-1}+2s_{i+n_E}+s_{i+n_E+1})-c-s_{i-1})=1/2(s_{i-n_E}+2s_{i-n_E-1}-2s_{i+n_E}-s_{i+n_E+1})-(c+1)$. This is positive only at the maximum, which is equal to $1/2-c$, and only if $c<1/2$. The maximum is attained at $\overline{s_i}=(a,a,E,E)$, hence $d\pi_i<0$ when $\overline{s_i}=(a,a,E,E)$, $c<1/2$, $(s_{i-1},s_{i+1})=(a,E)$ and $s_i=E$. □

Lemma 3.1 generalizes a result in ESS.[12] There it was shown that if $n_E=n_I=1$,

[12] The result is not formally stated in ESS, but appear on page 161.

then i chooses altruism if $(s_{i-2}, s_{i-1}, s_{i+1}, s_{i+2}) = (a, a, E, E)$. The reason is that i observes altruists that are grouped together and therefore have high payoffs. At the same time the egoists that he observes are also grouped together and therefore have low payoffs. When $n_E > 1$ what matter for i's decision are the actions of the individuals in N_i^I and just inside and just outside of N_i^E. The reason is that all the individuals in N_i^I *receive the same utility from the contributions of the individuals in* $\{i-n_E+1, ..., i-2, i+2, ..., i+n_E+1\}$. This is simply because $i-1$, i and $i+1$ share a large number of "externality neighbors". For example, suppose $n_E = 5$. In this case $i-1$, i and $i+1$ all receive the same utility from the actions of $i \pm 2, i \pm 3$ and $i \pm 4$. Therefore, the actions of these individuals are irrelevant for the computation of the *difference* between the average payoffs of altruists and egoists in N_i^I. However, the contribution of $i + 5$ is not received by $i - 1$, the contribution of $i - 5$ is not received by $i + 1$, and the contributions of $i + 6$ and $i - 6$ are only received by $i + 1$ and $i - 1$, respectively. The individuals in $\{i - 6, i - 5, i + 5, i + 6\}$ can be thought of as the "pivotal" contributors of N_i^I, i.e. the individuals whose actions benefit the individuals in N_i^I non-uniformly. Suppose $(s_{i-1}, s_i, s_{i+1}) = (a, s_i, E)$. In this case, for i to choose altruism, it is necessary that the actions of the pivotal contributors benefit the altruists but not the egoists in N_i^I. The ideal constellation, and indeed the only one such that i would choose altruism, is $(s_{i-6}, s_{i-5}, s_{i+5}, s_{i+6}) = (a, a, E, E)$. In this way, (the altruist) $i - 1$ but not (the egoist) $i + 1$ receive the contributions of $i - 6$ and $i - 5$. At the same time, $i + 5$ and $i + 6$ are egoists and therefore do not provide utility to $i + 1$, which would be out of reach of $i - 1$.

The next result characterizes the absorbing states of Γ. Let Ω_a and Ω_E denote the states in which all individuals are altruists and egoists respectively. We say that any $(s_{i-x}, ..., s_{i-2}, s_{i-1}, s_i, s_{i+1}, s_{i+2}, ..., s_{i+y})$ such that $s_j = a$ for all $j = i - x + 1, ..., i + y - 1$ and $s_{i-x} = s_{i+y} = E$ is a *string* of altruists of length $x + y - 1$. Correspondingly, any $(s_{i-x}, ..., s_{i-2}, s_{i-1}, s_i, s_{i+1}, s_{i+2}, ..., s_{i+y})$ such that $s_j = E$ for all $j = i - x + 1, ..., i + y - 1$and $s_{i-x} = s_{i+y} = a$ is a *string* of egoists of length $x + y - 1$.

Proposition 3.1. *Absorbing States of the Unperturbed Process*

Suppose $n_E > n_I = 1$.

(i) If $c \geq 1/2$, *then* $\Omega = \Omega_a \cup \Omega_E$.

(ii) With n *sufficiently large and* $c < 1/2$, *there are four categories of absorbing states:*

(1) Ω_a.

(2) Ω_E.

(3) *Strings of altruists of length at least* $n_E + 2$ *separated by strings of egoists of length* $n_E + 1$.

(4) *If* $n_E > 5$, *altruist strings of length between* 3 *and* $n_E - 3$ *separated by egoist strings of length between* 2 *and* $n_E - 4$.

(iii) A lower bound on the fraction of altruists in the absorbing states referred to in (3) is $\frac{n_E+2}{2n_E+3}$.

Proof. (i) This is a direct corollary of Lemma 3.1, since it shows that unless $c < 1/2$ an absorbing state cannot contain both altruists and egoists and evidently states containing either only altruists or only egoists are absorbing.

(ii) (1) and (2) are obvious.

(3): In an absorbing state Lemma 3.1 requires that in a straightened segment of the circle we have the following structure anywhere that an altruist is adjacent to an egoist[13]:

$$...a_{i-n_E-1}a_{i-n_E}........a_{i-1}a_iE_{i+1}........E_{i+n_E}E_{i+n_E+1}...$$

By further using Lemma 3.1 the following can be deduced: $[s_{i-1} = a \rightarrow s_{i-2} = a]$, $[s_{i-n_E} = a \rightarrow s_{i-n_E+1} = a]$, $[s_{i-n_E+1} = a \rightarrow s_{i+2} = E]$ and $[s_{i+1} = E \rightarrow s_{i+nI+2} = a]$. We then obtain:

$$...a_{i-n_E-1}a_{i-n_E}a_{i-n_E+1}\underset{\alpha}{......}a_{i-2}a_{i-1}a_iE_{i+1}E_{i+2}\underset{\beta}{.......}E_{i+n_E}E_{i+n_E+1}a_{i+n_E+2}... \quad \text{(A1)}$$

This structure must always result in an absorbing state wherever altruists and egoists meet and follows by iterating Lemma 3.1. Without this structure, some

[13] For the sake of concreteness the illustrations are made for the case $n_I = 10$, and each dot in the illustration represent an agent taking an unspecified action.

egoist would observe altruists better off than egoists, or some altruist would observe egoists better off than altruists. The question is how the empty spaces, denoted by α and β, can be filled. Write $\alpha = \{i - n_E + 2, \ldots, i - 3\}$, $\beta = \{i + 3, \ldots, i + n_E - 1\}$. These are of size $n_E - 4$ and $n_E - 3$ respectively. Note that A1 implies that egoist (altruist) strings must be of length 2 (3) at least and an egoist string can be at most of length $n_E + 1$. However, there is no corresponding upper bound to the length of altruist strings.

Next, we note that if α contains only altruists, which implies an altruist string of length at least $n_E + 2$, then β must contain only egoists. Otherwise any altruist in β with an egoist to his left will switch by Lemma 3.1. On the other hand, if β contains only egoists all individuals in α must be altruists. Otherwise an altruist in α with an egoist to his left will switch. This implies that if there is some altruist string of length at least $n_E + 2$ it must be bordered on each side by egoist strings of length $n_E + 1$ followed by more altruist strings of length at least $n_E + 2$. In the same way, if there is an egoist string of length $n_E + 1$ it must be bordered on each side by altruist strings of length at least $n_E + 2$ followed by egoist strings of length $n_E + 1$. Consequently in an absorbing state with an altruist string of length at least $n_E + 2$, all egoist strings are of length $n_E + 1$ and vice versa.

(4): Consider again A1

$$a_{i-n_E-1} a_{i-n_E} a_{i-n_E+1} \underbrace{\ldots\ldots}_{\alpha} a_{i-2} a_{i-1} a_i E_{i+1} E_{i+2} \underbrace{\ldots\ldots}_{\beta} E_{i+n_E} E_{i+n_E+1} a_{i+n_E+2} \tag{A1}$$

In order to have an egoist string of length less than $n_E + 1$ in an absorbing state there must be some altruists in β. We know that any altruist string must have length at least 3. Hence, the length of any egoist string shorter than $n_E + 1$ can be at most $n_E + 1 - 3 - 2 = n_E - 4$. Analogously, if there are some egoists in α there must be at least two of them, restricting the length of altruist strings to $n_E + 2 - 3 - 2 = n_E - 3$. If $n_E \leq 5$ then egoist strings would be at most of length 1 and by A1 we know that such a string cannot be part of an absorbing state. Therefore, with $n_E \leq 5$ all absorbing states are of the kind described in (3).

Naturally, any absorbing state with "short" strings as those discussed here must satisfy Lemma 3.1. The implication is that any egoist string must be part of a "broken" string of length $n_E + 1$ and altruists bordering egoists must have altruists at distance n_E and $n_E + 1$ on the side opposite of the egoist string. The following illustrates an absorbing state with $n = 18$ and $n_E = 10$:

$$\underbrace{EEaaaaaaaEE}_{\text{Broken egoist string}}aaaaaaa$$

There can never be short strings and long strings in an absorbing state, which follows by the proof of (3).

(iii) The lower bound of altruists in type 3 absorbing states is obtained simply by taking an altruist string of length $n_E + 2$ and an egoist string of length $n_E + 1$ and is therefore equal to $(n_E + 2)/(2n_E + 3)$. $\square$

Proposition 4.1 implies that a sufficient condition for the existence of absorbing states with both egoists and altruists is $n \geq (n_E + 2) + (n_E + 1) = 2n_E + 3$ or $n_E \leq (n - 3)/2$. This means that as long as the public good is not completely global, in the sense that there are at least two individuals in the population not enjoying the contribution of each individual, we can observe altruistic behavior, even in the presence of free riding egoists. Hence, altruists can coexist with egoists even when the externality is far less local than what is assumed in ESS. In fact, it is sufficient that the public good is non-global. On the other hand if the public good is completely global, there is no hope for altruism. This conclusion is emphasized in the following corollary:

Corollary 18. *There are absorbing states in which egoists and altruists coexist if* $n_E \leq (n - 3)/2$.

Proposition 4.1 also shows that several properties from the case $n_E = n_I = 1$ generalize naturally. The structure of the absorbing states in (3) of Proposition 4.1 is analogous to the case $n_E = n_I = 1$, in the sense that long strings of altruists are separated by strings of egoists. The intuition is similar too. By grouping together, altruists can exclude egoists from the benefits of cooperation. As imitation is local, this means that altruists observe mostly altruists with many altruist neighbors. The observed payoffs of altruists are therefore high, which leads them to be imitated. Similarly, egoists observe mostly egoists with many egoist neighbors. The observed payoffs of egoists are therefore relatively low and they are therefore not imitated. As the externality becomes less local, the groups of altruists must be larger for the contributions to be isolated from the egoists to a sufficient extent. Egoists live in groups on the edges of the altruist groups, benefitting to some degree from the public goods provided by these. However, these egoist groups can only grow to a certain limit, restricted by the size of the externality neighborhood. The lower bound for altruism derived here generalizes the lower bound derived by ESS for the case $n_E = 1$, which was found to be 0.6. Here, the lower bound is $\frac{n_E+2}{2n_E+3}$, which is equal to 0.6 at $n_E = 1$ and converges to 0.5 as n_E becomes large.

The absorbing states in (4) of Proposition 3.1 do not appear when $n_E = 1$. In these absorbing states, short altruists strings are mixed with short egoist strings. These absorbing states occur since altruists can benefit from distant altruist strings that the neighbor egoist do not benefit from. Absorbing states of type (4) can take a variety of forms. What they have in common is that the altruist and egoist strings are relatively short and that there are implicit "broken" egoist strings of length $n_E + 1$ as described in the proof.

3.4 Stochastically Stable States of The Perturbed Process

Even though absorbing states in which altruists and egoists coexist can be found when interaction is almost global, so far nothing has been said about the stability of these states in the presence of experimentation. Experimentation sometimes makes it more difficult for altruism to persist. For example, if $n_E = (n-3)/2$, then one experimentation is sufficient to lead an absorbing state in which there are both altruists and egoists into a state in which all individuals are egoists. On the other hand, a single experimentation is never enough to restore altruism. This means that if $n_E = (n-3)/2$, then egoism is somehow more stable than states in which there are both altruists and egoists. Even though altruism could appear and endure for some time even under these circumstances, it seems that egoism would have the upper hand in the long run. Nevertheless, if the population is larger, a greater number of experimentations are necessary to eliminate altruism. The reason is that in a larger population, there will either be more or longer altruist strings in an absorbing state, and there is a limit to the number of altruists that can be converted by a single egoistic experimentation. At the same time, the number of experimentations required to reintroduce altruism into a world of egoism is never above $n_E + 1$. Hence, the stability of altruism in the presence of experimentation is related to both n_E and n. To analyze this situation formally, this section uses techniques developed by Freidlin and Wentzell (1984) and introduced into economics by Kandori, Mailath and Rob (1993) and Young (1993) in order to characterize the stochastically stable states of Γ^ε. As mentioned, the process will spend almost all of its time in the stochastically stable states in the long run. The characterization of stochastic stability is done for the case $n_E \leq 5$. The reason is that when $n_E > 5$ absorbing states of the type in (4) of Proposition 3.1 emerge and it is particularly complicated to characterize the basin of attraction of these states.

Let Ω_3 denote the absorbing states specified in (3) of Proposition 3.1. Let $\lceil x \rceil$ be the smallest integer greater than x.

Proposition 3.2. *Suppose* $1 = n_I < n_E \leq 5$. *(1) If* $n > 4(n_E+1)^2$, *then* $\Omega^* = \Omega_3$. *If* $\left\lceil \frac{n}{4(n_E+1)} \right\rceil > n_E + 1$, *then* $\Omega^* = \Omega_E$. *If* $\left\lceil \frac{n}{4(n_E+1)} \right\rceil = n_E + 1$, *then* $\Omega^* = \Omega_E \cup \Omega_3$.

Proof. The methodology of Young (1993, 1998) is used to characterize Ω^*. In short, the method works as follows: For any $\omega \in \Omega$, an ω-tree is a tree branching out from ω and reaching every $\omega' \in \Omega$, $\omega' \neq \omega$, through a unique path. The cost $c(\omega', \omega'')$ of the edge between ω' and ω'' in this graph is equal to the minimum number of experimentations required to move from ω' to ω''. Let Θ_ω be the set of all ω-trees.

The stochastic potential of ω is defined as $\gamma(\omega) = \min\limits_{\theta \in \Theta_\omega} \sum_{(\omega', \omega'') \in \theta} c(\omega', \omega'')$. Then $\omega \in \Omega^*$ if and only if it is $\omega \in \arg\min\limits_{\omega' \in \Omega} \{\gamma(\omega')\}$. In other words, Ω^* corresponds to the states with minimum stochastic potential.

The proof is carried out in three steps: (1) It is shown that all states in Ω_3 have the same stochastic potential, which implies that if some $\omega \in \Omega_3$ is stochastically stable then all $\omega' \in \Omega_3$ are stochastically stable. (2) It is shown that Ω_a is never stochastically stable. (3) The conditions under which either Ω_a or Ω_E are stochastically stable are derived.

Step 1: *Any state $\omega \in \Omega_3$ can be reached by a series of 1-experimentation transitions from any other state in Ω_3.*

By a "x-experimentation transition" we mean that x experimentations combined with imitation dynamics lead the system from one absorbing state to another. This is abbreviated 1ET and $\omega \xrightarrow{x} \omega'$ denotes such a transition between ω and ω'. Let $x(k)$ denote the set of type 3 absorbing states in which there are $k \geq 1$ strings of egoists. We start by making two observations:

(i) A single experimentation can destroy an altruist string in a state $\omega \in \Omega_3$ of length at most $2(n_E + 2) + (n_E + 1) - 2 = 3n_E + 3$. We call such a string a *short string* and a longer altruist string a *long string.* The destruction of a short string takes place when an individual in its middle experiments to egoism and subsequently converts the entire altruist string step by step. If there is some other altruist string, there is now a new long string of egoists that will shrink at its edges until it reaches length $n_E + 1$, in this way making the adjacent altruist strings grow. The presence of a short string thus makes a 1ET from a state in $x(k)$ to a state in $x(k-1)$ possible. It is important to note that with positive probability, only one of the adjacent altruist strings grows after the destruction of the altruist string.

(ii) A long altruist string can be divided into two strings by a single experimentation which introduces an egoist into its middle, which grows to length $n_E + 1$, making possible a 1ET from a state in $x(k)$ to a state in $x(k+1)$.

We now provide an algorithm in two steps for moving from any state in $x(k)$ to a state in $x(1)$ through a sequence of 1ET. If we denote the set of states in which there are k_1 short strings and k_2 long strings by $x(k_1, k_2)$, then we finish when $k_i = 0$ and $k_j = 1$, $i \neq j$.

(i) Here the goal is to show that from any state in $x(k_1, k_2)$ a state in $x(k_1, 1)$ can be reached through a series of 1ET. If $k_2 = 1$ from the beginning, go to (ii). If $k_2 > 1$, fragmentize all long strings but one by putting experimentations to egoism at appropriate places one at a time and allowing each new egoist strings to grow to size $n_E + 1$. We end up in a state in $x(k_1 + z, 1)$, $z > 0$. If $k_2 = 0$ destroy short strings and let an adjacent altruist string grow until a long string is created and we are in $x(k_1 - m, 1)$ or only one short string remains and we are in $x(1, 0)$, in which case we are finished.

(ii) We will now be in a state in $x(k_3, 1)$, where $k_3 \in \{k_1 + x, k_1, k_1 - m\}$. Start merging strings with the long string. This is done by suitably placing a experimentation inside a short altruist string adjacent to the long altruist string and constructing a sequence of imitations such that the long egoist string created in this way is completely absorbed into the long altruist string. Proceeding in this way with each of the short strings, they are eventually eliminated and we arrive in $x(0, 1)$, which was the goal.

Next, any state in $x(1)$ can be reached from any other state in $x(1)$ by a series of 1ET ($x(1)$ contains more than one state since the egoist string may be at different locations). To "move" the egoist string, let an altruist adjacent to the egoist string mutate and next let the egoist at the opposite end of the egoist string imitate to altruism. In this way, one experimentation moves the egoist string one step. Thus any state in $x(1)$ can be reached by 1ETs from any other state in $x(1)$.

Finally, any state in $x(k)$, $k > 1$, can be reached from a state in $x(1)$ via a series of 1ETs. To accomplish this, for any state $\omega \in x(k)$, pick a state in $x(1)$ with a string of egoists that coincide in location with some egoist string in ω and distribute experimentations in the centre of the location of all egoist strings in ω in any order (letting the imitation dynamics work after each experimentation). The system will converge to ω via a series of 1ETs.

Consequently any state $\omega \in \Omega_3$ can be reached from any other state in Ω_3 via a series of 1ETs. This implies that all states in Ω_3 have the same stochastic potential.

Step 2:

First note that $\omega_A \xrightarrow{1} \omega$ is possible for some state $\omega \in \Omega_3$. Consider the ω_A-tree. In the path from ω to ω_A at some point there is an edge leaving some state ω' at cost greater than 1, since egoist strings are always at least of length 2 in absorbing states and 2 experimentations therefore are necessary to eliminate egoism. Cut that edge and add the edge $\omega_A \xrightarrow{1} \omega$ and we get a ω'-tree with lower total cost than the ω_A-tree. Hence, ω_A is not stochastically stable.

Step 3:

Consider first the ω-tree for $\omega \in \Omega_3$. In such a tree the edges exiting any $\omega' \in \Omega_3$ as well as Ω_a must have cost 1. Next, Ω_E is connected to the ω-tree in the least costly way. It can be shown that only altruist strings at least of length $n_E + 1$ can grow to size $n_E + 2$ or larger. At the same time, the introduction of a $n_E + 1$ string of altruist is sufficient for the system to reach a state in Ω_3. Hence, in the ω-tree the edge leaving Ω_E has cost $n_E + 1$.

Now consider the Ω_E-tree. In the Ω_E-tree the edge leaving Ω_a has cost 1 and is connected to some state in Ω_3. All states in Ω_3 have outgoing edges of cost 1, except one that is connected to Ω_E. This is the state $\omega \in \Omega_3$ from which Ω_E can be reached at the lowest cost. A single experimentation can destroy an altruist string of length at most $3(n_E + 1)$. The largest population in a type (3) absorbing state that can be converted to egoism by m experimentations is $\hat{n} = m((n_E + 1) + 3(n_E + 1)) = 4m(n_E + 1)$. This occurs in the state in which altruist strings are of length precisely $3(n_E + 1)$. Consequently $m^* = \left\lceil \frac{n}{4(n_E+1)} \right\rceil$ is the smallest number of experimentations required to move from a state in Ω_3 to Ω_E.

Whether Ω_3 or Ω_E is stochastically stable then depends on whether $n_E + 1$ is larger than $\left\lceil \frac{n}{4(n_E+1)} \right\rceil$. Then Ω_3 is stochastically stable if $\left\lceil \frac{n}{4(n_E+1)} \right\rceil > n_E + 1 \rightarrow \frac{n}{4(n_E+1)} > n_E + 1$ or $n > 4(n_E + 1)^2$. If $\left\lceil \frac{n}{4(n_E+1)} \right\rceil < n_E + 1$ holds Ω_E is stochastically stable and in case of equality both are stochastically stable. $\square$

The first conclusion from Proposition 4.2 is that the absorbing state in which all agents are altruist is not stochastically stable and will therefore not appear in the long run. The reason is that it is very sensitive to experimentation. A single egoistic experimentation thrives in this case and grows into string of length $n_E + 1$ with positive probability. In this way we a reach a state in which there are both altruists and egoists. At the same time, it is hard to reach a state where all agents are altruists from another absorbing state, basically because this requires simultaneous experimentation of many egoists. It is more difficult to leave states in which there are both altruists and egoists for a state where all individuals are either altruists and egoists. Proposition 4.2 shows that such states are sustainable when the population is sufficiently large. Moreover, the required size of the population increases with n_E. Indeed, altruistic behavior is possible for any size of the externality neighborhood, if the population is correspondingly large. The required size of the population increases quadratically with n_E. The intuition is that as n_E increases, altruist strings must be longer to persist and grow. More experimentations are then needed to introduce altruism into a world of egoism. At the same time, fewer experimentations are required to eliminate altruism and descend into egoism. In this sense, for a fixed population size, a large externality neighborhood is bad for altruism and eventually precludes altruism in the long run. However, the number of experimentations required to exterminate altruism increases linearly with population size. As the population grows, at some point it becomes so difficult to eradicate altruism that it ends up prevailing in the long run. A large population protects altruism in the sense that it becomes increasingly difficult for egoistic experimentations to eliminate all altruists at once. If some agent switches from altruism and causes a temporary increase in egoistic behavior, there are pockets of altruism at other population that are unaffected by this burst of egoism.

Proposition 4.2 generalizes the conclusions of ESS, who only show that when $n_E = 1$, the required size of the population is 30. Proposition 4.2 allows us to make a precise statement with regard to the required localness of the public good for altruism to persist in the presence of experimentations. A measure of the globalness of the public good is the fraction $2n_E/n$, i.e. the fraction of the population enjoying the contribution of each individual. From Proposition 4.2 we have that $n > 4(n_E + 1)^2$ is necessary and sufficient for all stochastically stable states to involve altruism. This expression can be rewritten $\frac{n_E}{2(n_E+1)^2} > \frac{2n_E}{n}$, which means that n_E puts an upper bound to the globalness of the public good. The term $\frac{n_E}{2(n_E+1)^2}$ is maximized at $n_E = 1$ at which $n \geq 17$ is necessary for altruism to prevail.[14] Hence, at most 2/17 individuals of the population can enjoy the contributions of each individual for all stochastically stable states to involve altruism. In other words, the public good must be quite local in the relative sense considered here for altruism to prevail in the long run. This finding is summarized in the following corollary:

Corollary 19. *For altruism to prevail in the long run, the externality can benefit at most a fraction 2/17 of the population. This occurs when* $n_E = 1$*.*

The intuition behind the result is that with larger n_E, altruist strings must be larger in order to exclude egoists from the contributions to a sufficient degree. If the population is small, not many such strings fit into the population, which makes absorbing states sensitive to egoistic experimentation. If the population is large, many simultaneous egoistic experimentations are needed to eliminate altruism. This is the case since many altruistic strings can fit into a large population and a small number of egoistic experimentations would leave pockets of altruists which subsequently thrive. However, it turns out that the required population size increases fast, and indeed altruism can be stochastically stable if the externality reaches at most a fraction 2/17 of the population.

3.5 Larger Imitation Neighborhoods

This section presents a result for the case in which n_I can be greater than one, but it still holds that $n_I \leq n_E$. The model is more complex in this case and it is difficult to provide a complete characterization of the outcome, even more so since the dynamics

[14]The difference with respect to ESS comes from the fact that ESS considers $\mu = 1$, i.e. all individuals necessarily revise their choice in each period, whereas here $\mu \in (0, 1)$.

depend to a considerable extent on c. Instead, I characterize a structure of intuitive appeal, composed of long strings of altruists separated by long strings of egoists, that will constitute an absorbing state given that c is in a certain range and given that n_E is significantly larger than n_I. This structure also encompasses the type 3 absorbing states characterized in Proposition 3.1.

Proposition 3.3. *With n sufficiently large and $n_E \geq 2n_I - 1$, strings of altruists of length at least $2n_I + n_E$ separated by egoist strings of length $n_I + n_E$ are absorbing states of the unperturbed process, if and only if $c \in [\frac{2(n_I)^2+n_I-3}{2(1+n_I)}, \frac{2n_I-1}{2})$.*

Proof. The following illustration of a typical absorbing state (with $n_I = 3$ and $n_E = 7$) is provided as a point of reference:

$$aaaaaaaaaaaa\underset{r}{a}\underset{s}{E}EEEEEEEEE \tag{3.1}$$

Let r and s be the altruistic and egoistic individuals between two strings, as in the illustration. Index $r = s = 0$ and index the individuals to the left of r according to position relative to r, $\{-(r-n_I), (r-n_I+1), \ldots, r\} = \{n_I, n_I-1 \ldots, 1, 0\}$, and to the right of s according to $\{v, v+1, \ldots, v+n_I\} = \{0, 1, \ldots, n_I\}$. Let $\{-(r-n_I), (r-n_I+1), \ldots, r\} \equiv A$ and $\{v, v+1, \ldots, v+n_I\} \equiv B$. A and B then correspond to the altruists with some egoist in their imitation neighborhood and the egoists with some altruist in their imitation neighborhood, respectively.

We first consider altruists. Denote the highest cost at which an altruist remains so by $\bar{c}$. This means that we should find $\bar{c}$ such that $d\pi_i(\bar{c}) > 0$ for all $i \in A$ and for all $c \in [0, \bar{c})$. For an altruist string of length $2n_I + n_E$:

$$\begin{aligned}
d\pi_i(c) &= \frac{1}{n_I+1+i}\sum_{k=0}^{n_I+i}(n_E+k) - \frac{1}{n_I-i}\sum_{k=0}^{n_I-i-1}(n_E-k) - c \\
&= \frac{1}{n_I+1+i}\sum_{k=0}^{n_I+i}(k) - \frac{1}{n_L-i}\sum_{k=0}^{n_I-i-1}(k) - c \\
&= \frac{(n_I+1+i)(n_I+i)}{2(n_I+1+i)} - \frac{(n_I-i)(n_I-i-1)}{2(n_L-i)} - c \\
&= \frac{2n_I-1}{2} - c.
\end{aligned}$$

The first and second term are the average payoffs of altruists and egoists in N_i^I respectively. Note that we can write this only if $n_E \geq 2n_E - 1$. We thus have that $\bar{c} = \frac{2n_I - 1}{2}$, which is independent both of i (the position of the individual), and n_E. If the altruist string is longer than $2n_I + n_E$, the only difference is that $d\pi_i(c)$ becomes larger for some of the individuals in the interior of the altruist string and $\bar{c}$ will thus still provide the highest cost at which altruists remain so. Hence $c < \bar{c}$ is necessary and sufficient for altruists not to switch in a state in which strings of altruists of length at least $2n_I + n_E$ are separated by egoist strings of length $n_I + n_E$.

We now turn to the case of the egoists. We want to find $\underline{c}$ such that $d\pi_i(c) \leq 0$ for all $j \in B$ and for all $c > \underline{c}$, with the additional requirement that $\underline{c} < \bar{c}$. This part of the proof is carried out in two steps: (1) we find $\underline{c}$ by solving $d\pi_s(c) = 0$ and then check that $\underline{c} < \bar{c}$. (2) We show that $d\pi_j > d\pi_{j+1}$, for $j, j+1 \in B$ which ensures that no individual in B wants to switch.

(1) We can write

$$\begin{aligned} d\pi_s(c) &= \frac{1}{n_I} \sum_{k=0}^{n_I - 1} (n_E + k) - \frac{1}{n_I + 1} \sum_{k=0}^{n_I - 1} (n_E - k) - c \\ &= \frac{2(n_I)^2 + n_I - 3}{2(1 + n_I)} - c, \end{aligned}$$

which means that $\underline{c} = \frac{2(n_I)^2 + n_I - 3}{2(1+n_I)} \geq 0$. The egoist s thus remains an egoist if and only if $c \geq \underline{c}$. Next, $\bar{c} - \underline{c} = \frac{2n_I - 1}{2} - \frac{2(n_I)^2 + n_I - 3}{2(1+n_I)} = \frac{1}{1+n_I} > 0$.

(2) We now show that $d\pi_j > d\pi_{j+1}$, for $j, j+1 \in B$ and show that this is greater than 0. We can write $d\pi_j - d\pi_{j+1} = \bar{\pi}_j^a - \bar{\pi}_{j+1}^a - (\bar{\pi}_j^E - \bar{\pi}_{j+1}^E)$, $\bar{\pi}_j^{a/E}$ denote the average payoffs of altruists and egoists respectively in N_j^I. Then

$$\bar{\pi}_j^a - \bar{\pi}_{j+1}^a = \frac{1}{n_I - j} \sum_{k=0}^{n_I - j - 1} (n_E + k) - \frac{1}{n_I - j - 1} \sum_{k=0}^{n_I - j - 2} (n_E - k) = \frac{1}{2}.$$

Next by taking into account the overlap of N_j^I and N_{j+1}^I:

$$\begin{aligned}
\bar{\pi}_j^E - \bar{\pi}_{j+1}^E &= \frac{1}{n_L + j + 1} \sum_{k \in N_j^I : s_k = E} \pi_k - \frac{1}{n_I + j + 2} \sum_{k \in N_j^I : s_k = E} \pi_k - \frac{\pi_{j+n_I+1}}{n_I + j + 2} \\
&= \frac{\sum_{k \in N_j^I : s_k = E} \pi_k - (n_I + j + 1)\pi_{j+n_I+1}}{(n_I + j + 1)(n_I + j + 2)} = \frac{\bar{\pi}_j^E - \pi_{j+n_I+1}}{n_I + j + 2}
\end{aligned}$$

The following bound can be established:

$$\begin{aligned}
\bar{\pi}_j^E &\leq \frac{\sum_{k=0}^{n_I - 1} n_E - k}{n_I} = n_E - \frac{n_I - 1}{2} \\
\pi_{j+n_I+1} &\geq n_E - (n_I - 1)
\end{aligned}$$

$\bar{\pi}_j^E$ is bounded by $\bar{\pi}_s^E$ and π_{j+n_I+1} is bounded by the smallest payoff in the egoist string. We can then write:

$$\bar{\pi}_j^E - \bar{\pi}_{j+1}^E = \frac{\bar{\pi}_j^E - \pi_{j+n_I+1}}{n_I + j + 2} \leq \frac{1}{2} \frac{n_I - 1}{(n_I + j + 2)} < \frac{1}{2} = \bar{\pi}_j^a - \bar{\pi}_{j+1}^a$$

Consequently, $d\pi_j > d\pi_{j+1}$. Hence $c \geq \underline{c}$ is necessary and sufficient for egoists not to switch in a state in which strings of altruists of length at least $2n_I + n_E$ are separated by egoist strings of length $n_I + n_E$. □

Proposition 4.3 shows that absorbing states in which long strings of altruists are separated by long strings of egoists exist also when imitation neighborhoods are larger. When $n_I = 1$, the case analyzed in the previous sections is obtained. The condition $n_E \geq 2n_I - 1$ implies that externality neighborhoods that are much larger than the imitation neighborhoods are considered. The upper bound for the cost is the highest cost at which altruists remain so. This increases with n_I. The intuition is that as imitation neighborhoods become larger, altruists can see deeper into both altruist and egoist strings where altruists are happier and egoists are more miserable. Hence, altruists remain so at larger costs. The lower bound for the cost is the lowest cost at which egoists remain so. This also increases with n_I and it does so more rapidly than the upper bound. Hence, the interval becomes smaller as n_I becomes larger. However, this is because egoists become more tempted to become altruists and not the opposite. At costs below the lower bound it should be even easier for altruism to thrive. Therefore, since larger n_I increases the upper bound for the cost, more information in a sense helps altruism. This should be contrasted with the conclusion in Mengel (2009), that too large imitation neighborhoods are detrimental for altruism. When the imitation neighborhoods are smaller than the externality neighborhoods, more information actually seems to be beneficial for altruism.

3.6 Concluding Remarks

This paper has analyzed the importance of localness for altruism in a model of local provision of a public good. The conclusion is that in the short run, which is reflected by the process without experimentations, altruism can survive as long as the public good is not global. With experimentation the process tends to select against altruism in the long run when the public good is too global (in the sense of reaching a larger fraction of the population). In fact, altruism will not survive in the long run if the public good is more global than in the case analyzed by ESS. As a larger set of neighbors enjoy the contribution of each individual, the required size of the population for altruism to survive in the long run increases fast.

The model considered in this paper is to stylized to have policy implications. However, it does suggest that if one wants to enforce altruistic behavior, it may be better to concentrate enforcement locally. For example, it may be necessary to enforce emission control of automobiles. An implication of this paper is that to the extent that the benefits of implementing emission control are sufficiently local, it may be better to concentrate the efforts of enforcement to certain areas, instead of spreading them over the entire population. In this way, the benefits of reducing emissions can be revealed and subsequently imitated by other segments of the population. It may even be a good idea to temporally enforce altruism locally, even if this investment has a negative benefit net cost in the short run. If behavior is driven by imitation, the locally enforced altruism can spread and eventually the enforcement may not be necessary. However, care should be taken since this can work only if the population is much larger than the reach of the externality. Otherwise, in the long run, it is likely that the population is driven back to a state of egoistic behavior.

Bibliography

[1] Apesteguia J., Huck, S. and Oechssler, J. (2007). Imitation - Theory and Experimental Evidence. *Journal of Economic Theory 136*, 217-235.

[2] Becker, G.S. (1974). A theory of social interactions. *Journal of Political Economy 82*, 1063-1093.

[3] Becker, G.S. (1981). A Treatise on the Family. Cambridge: Harvard University Press.

[4] Binmore, K., Samuelson, L. and Vaughan, R. (1995). Musical Chairs: Modeling Noisy Evolution. *Games and Economic Behavior 11*, 1-35.

[5] Eshel I., Samuelson L. and Shaked A. (1998). Altruists, Egoists and Hooligans in a Local Interaction Model. *American Economic Review 88*, 157-179.

[6] Freidlin, M. and Wentzell, A.D. (1984). Random Perturbations of Dynamical Systems. New-York: Springer-Verlag.

[7] Fudenberg, D. and Maskin, E. (1986). The Folk Theorem in Infinitely Repeated Games with Discounting or with Incomplete Information. *Econometrica 54*, 533-554.

[8] Huck, S., Normann, H.T, and Oechssler, J. (1999). Learning in Cournot Oligopoly: an Experiment. *Economic Journal 109*, 80-95.

[9] Jun, T. and Sethi, R. (2007). *Journal of Evolutionary Economics 17,* 623-646.

[10] Kandori, M., Mailath, G. and Rob, R. (1993). Learning, Mutation and Long Run Equilibria in Games. *Econometrica 61*, 29-56.

[11] Karlin, S. and Taylor, HM. (1975). A First Course on Stochastic Processes. Academic, San Diego.

[12] Kirchkamp, O. and Nagel, R. (2007). Naïve Learning and Cooperation in Network Experiments. *Games and Economic Behavior 58*, 269-292.

[13] Matros, A. 2008. Altruistic Versus Rational Behavior in a Public Good Game. University of Pittsburgh, Working Paper.

[14] Mengel, F. 2009. Conformism and Cooperation in a Local Interaction Model. *Journal of Evolutionary Economics 19*, 397-415.

[15] Selten, R. and Apesteguia J. 2005. Experimentally Observed Imitation and Cooperation in Price Competition on the Circle. *Games and Economic Behavior 51*, 171-192.

[16] Young, P. 1993. The Evolution of Conventions. *Econometrica 61*, 57-84.

Chapter 4

Perfect Communication with Arbitrary Communication Costs

4.1 Introduction

Strategic communication is an important aspect of many economic situations in which some party is privately informed. The so called *Persuasion Games* introduced by Paul Milgrom in 1981 in an influential paper, focus on strategic communication in terms of disclosure of verifiable or "hard" information. In the benchmark model, a privately informed *sender* aims to influence the behavior of a *receiver* by verifiably communicating information in a written report. The sender decides which information to withhold and which information to include in the report. Milgrom (1981) showed that if communication is costless and preferences satisfy a monotonicity requirement, the sender reveals all his private information to the receiver. In other words, communication is perfect in the presence of both incongruent preferences and strategic behavior. This frequently cited result is known as the *unraveling result.*[1]

However, whereas Milgrom (1981) assumed communication to be costless, reporting verifiable information is often costly. First, it often requires careful and detailed explaining based on facts. For example, consider an entrepreneur who writes a business plan to convince a venture capitalist to invest in his business. It takes

[1] Different generalizations of the result originally derived by Milgrom in 1981 can be found in Milgrom and Roberts (1986), Seidman and Winter (1991) and Mathis (2008).

both time and effort to explain the technical details of the product, the production costs, existing and possible competitors, the financial situation and other relevant information. Second, it may involve costly certification by accredited institutions. The financial information in the business plan perhaps requires auditing by external accountants, or a patent may be needed to certify the product's originality. At the same time, the existing literature indicates that unraveling is sensitive to the assumption of costless reporting. Jovanovic (1982), Verecchia (1983) and Cheong and Kim (2004) note that if reporting is costly, it becomes too expensive to report all available information and therefore unraveling cannot occur.

This paper delivers a different conclusion. I study a model based on Milgrom's (1981) benchmark Persuasion Game, in which, in contrast to previous work, the cost of verifiably reporting private information continuously increases in the precision of the report. The conclusion is, contrary to the previous literature, that unraveling is rather robust to costly reporting. More specifically, whereas Milgrom (1981) showed that if reporting is costless the unique equilibrium is separating, I show that a separating equilibrium always exists, regardless of the reporting costs. Hence, communication can be perfect even with arbitrarily high communication costs. The intuition is that the reporting costs introduce costly signaling into the Persuasion Game and thereby give the report a double function. It both discloses information and at the same time functions as a signaling device through the costs incurred producing it. When it becomes too expensive to report all the information, a high sender type can instead discourage lower types from mimicking his report through the reporting costs. It turns out that a combination of information disclosure and costly signaling always can accomplish full separation.

I further show that there may be several different separating equilibria, but these are all payoff equivalent. Moreover, when the reporting costs are low enough, all equilibria are separating. The uniqueness of the separating equilibria arises since with sufficiently low reporting costs, high sender types can always break out of any pooling equilibrium by disclosing their true type. When the costs instead are high, a pooling equilibrium emerges. In the pooling equilibrium, the reports contain no information at all. The intuition is that when the costs are high, skepticism on part of the receiver with respect to withheld information makes it too expensive for the sender to disclose his true type and break out of the pooling equilibrium. Hence, whereas the existence of a separating equilibrium is not sensitive to reporting costs, meaning that unraveling can occur also with high costs, uniqueness indeed is.

The reason that Jovanovic (1982), Verecchia (1983) and Cheong and Kim (2004) reach such a different conclusion, is that they assume that the sender's decision is binary. He must either report all his private information, or none of it. In such a setting, there are always some sender types that will not disclose their information and if the cost is too high, the sender never reports any information.[2] The assumption that the sender's decision is binary, means that he cannot send intermediate amounts of information. Therefore, if the costs are high, he cannot adapt his reporting strategy accordingly and report less information. Here, the sender instead has considerable discretion in this respect. More specifically, as in Milgrom's (1981) original Persuasion Game, the sender continuously decides how much information to report. This means that even if the reporting costs are arbitrarily high, the sender can report an arbitrarily small amount of information without incurring too high costs. The assumption that the sender's decision is continuous seems well suited to situations in which there is significant discretion with respect to how much information to include in the report, such as in the example of the entrepreneur writing a business plan.

Finally, an extension of the benchmark Persuasion Game is introduced in which the receiver has a more active role in the communication. In this extension, the receiver must make an effort at a cost in order to access the information contained in the report. The aim is to capture the fact that just as it is costly to elaborate a report, it frequently requires both time and effort to properly read and understand it.

Once the receiver actively decides whether to read a report, she may be able to

[2] The latter conclusion is also obtained by Eso and Galambos (2008), who also consider binary disclosure decisions.

condition this decision on her "first impression" of it. It then becomes important how this first impression is related to the information content of the report and to which extent the sender can manipulate it. Here, I assume that the appearance of a report is related to the amount of information it contains and that appearances can be manipulated at a cost.

Two classes of separating equilibria are characterized. In one of these, referred to as *non-reading equilibria*, the receiver never reads any report. Depending on the reporting and manipulation costs, the sender either refrains from manipulating the appearance of the report, or spends resources to make the report look more precise than what it is. In the other class of equilibria, referred to as *reading equilibria*, the receiver reads all the reports. Since both equilibria are separating, the receiver prefers the former, in which she incurs no effort costs.

The structure of the paper is as follows. Section 2 discusses some related literature. In Section 3 the model is introduced. Section 4 characterizes the equilibria of the benchmark model of costly reporting and discusses their properties. In Section 5 the extension of the benchmark model is presented. Section 6 concludes.

4.2 Related Literature

This paper is related to Mathis (2008), which considers communication in terms of *partially* verifiable information. Mathis (2008) postulates that sometimes due to time or technical constraints, it is impossible to verifiably report all private information. His conclusion, if applied to a standard Persuasion Game (Milgrom, 1981), is that the sender must be able to verifiably report all favorable information for a separating equilibrium to exist. Mathis' (2008) approach to partial verifiability can be treated as a model of costly reporting in which reports are either costless or arbitrarily costly. In other words, the reporting costs are discontinuous. The present paper gives an alternative approach to partial verifiability. Here, it may also be arbitrarily costly, and hence in practice impossible, to report all private information. However, in contrast to Mathis (2008), here a separating equilibrium always exist.

The difference arises due to the continuity of the reporting costs in the present paper. This allows the sender to use the reporting costs as a signaling device when it becomes too expensive to report all the information, which is not possible in Mathis (2008). Hence, when treating partial verifiability in terms of costly reporting, the continuity of the costs has an important impact on the outcome.

Another related paper is Kartik's (2009) work on costly lying.[3] In Kartik's model, the sender can provide false information at a cost. However, in contrast to the present paper, the cost of the report is unrelated to its precision. In a sense, Kartik's paper and this paper consider two different kinds of costly lying. Whereas in Kartik (2009) it is costly to provide false information, here the costs are related to how much relevant information is withheld. At the same time, while in Kartik (2009) lying is costly, here it is being truthful that is costly, in the sense that more precise reports cost more. Kartik (2009) finds that full separation is impossible and instead characterizes equilibria in which low types separate and high types pool. There is still no paper that accomodates both of the mentioned types of costly lying.

Note that the idea here is different from the one in Henry (2009) and in the section "Pecuniary Externalities of Disclosure" in Milgrom (2008). In these papers the sender chooses a number of costly tests to carry out. He then decides which tests to disclose to the receiver. In principle, this means that the sender chooses the precision of his private information as well as the precision of the reports. However, it is not costly to *report* the information. Hence, the focus of these models is on costly acquisition of information rather than on costly reporting of information, as in the present paper. The results are also different. Since reporting is costless in Henry (2009), the unraveling result holds and the sender always reveals all of his private information, which is not the case here.[4]

The extension in which the receiver must exert effort to assimilate the infor-

[3]It is also related to Kartik, Ottaviani and Squintiani (2006), but not as closely. The major difference is that Kartik et. al. (2006) consider an unbounded typespace, whereas Kartik's (2009) typespace is bounded, which is also the case in the present paper.

[4]By the same argument this paper is different from that of Austen-Smith (1994), in which the acquisition of information is costly.

mation in the report is related to Dewatripoint and Tirole (2005). In their model, the receiver also incurs a cost in order to understand reports, but in a significantly different setting. For example, Dewatripoint and Tirole (2005) consider a binary type space and a binary outcome of communication.

4.3 The Model

A game of persuasion Γ is considered. There are two players, a sender S and a receiver R. The game proceeds in two stages. In the first stage nature reveals the value of a parameter $t \in T = [0,1]$ to S, according to probabilities given by some known density f with full support. The parameter t is referred to as the sender's type and the knowledge of t constitutes his private information. The sender then delivers a report to the receiver, that takes the form of a closed interval contained in T. He can send any such interval with the restriction that his type must belong to it. Formally, let $M = \{[l,h] : 0 \leq l \leq h \leq 1]\}$ and let $M(t) = \{m \in M : t \in m\}$. A sender of type t then chooses a report $m \in M(t)$, which he delivers to the receiver[5]. Let the *precision* of a report $[l,h]$ be a function $\upsilon : M \to [0,1]$, defined as $\upsilon(l,h) = h - l$.[6]

This configuration allows the sender to be as vague or precise as he wants in his report. For example, $[0,1]$ is uninformative about the sender's type, whereas $[t,t]$ completely identifies him. The sender thus communicates an upper and lower bound for his type, and is restricted to do so truthfully. This captures the assumption that information is verifiable and the sender cannot include false information in the report, but is free to choose which information to transmit and which to withhold.

Reports are costly to produce. I assume that the cost of producing a report depends only on its precision. Hence, the cost is a function $C : [0,1] \to \mathbb{R}$. The cost does not depend on the sender's type, which is a significant difference from many

[5] I follow convention and let m, which usually stands for *message* in the literature on strategic transmission of information, denote an arbitrary report.

[6] The function υ is defined on closed intervals, so the correct notation would be $\upsilon([l,h])$, but I use the convention of eliminating the brackets in such expressions to lighten notation.

models of costly signaling. For example, it is equally costly for any type in $[l, 1]$ to send $[l, 1)$.[7]

After having received a report m, the receiver forms a posterior belief $\mu(t|m)$ with respect to the type that sent the report. This belief is a probability distribution over T. The beliefs are said to be *skeptical* with respect to $[l, h]$ if $\mu(t|[l, h])$ is degenerate at l. After forming her beliefs, the receiver chooses an action $a \in A$, where A is a closed interval.

The sender's payoff is given by a function $U : A \times M \to \mathbb{R}$, defined by $U(a, m) = u^S(a) - kC(v(m))$. The term $u^S(a)$ captures the dependence of the sender's payoff on the receiver's action and $k > 0$ parameterizes the intensity of the reporting cost. The receiver's payoff is given by a function $u^R : A \times T \to \mathbb{R}$. Hence, the receiver's payoff depends on both the action she chooses and the sender's type. For example, the sender's type may be the quality of a product that he offers for sale and the receiver's action the quantity that she decides to buy. In this case, the payoff of the receiver depends on both the quality of the product and the amount she decides to buy. The payoff of the sender depends only on the amount that he persuades the receiver to buy, and the costs of his persuasive efforts, i.e. of the report. Let $a^R(t) = \arg\max_{a \in A}\{u^R(a, t)\}$, i.e. $a^R(t)$ is the optimal action of the receiver given the type of the sender. The following is assumed with respect to payoff and cost functions

Assumption 1: C is continuous and decreasing.

Assumption 2: u^S is continuous and strictly increasing.

Assumption 3: For all $t \in T$, $a^R(t)$ is unique, continuous and strictly increasing.

Assumption 1 implies that more precise reports are more expensive to produce. Assumption 2 means that the sender's payoff increases in the receiver's action, which

[7] There are applications when this assumption is not the most natural one. For example, a researcher with a relatively good model may find it easier than a researcher with a bad one, to prove to colleagues that his model is not of the worst possible kind.

is natural if for example the action of the receiver is a purchased quantity or a level of investment. Assumption 3 implies that the receiver's optimal choice continuously increases in the type of the sender. This is reasonable if for instance the type represents the quality of the sender's product. Assumption 2 and 3 are the kind of monotonicity assumptions used by Milgrom (1981) and Milgrom and Roberts (1986) to derive the unraveling result. The essence of these two assumptions is that the sender wants the receiver to believe he is of as high a type as possible.

A pure strategy of the sender is a function $m : T \to M$ with the constraint $m(t) \in M(t)$ for all $t \in T$. When it is convenient to be explicit about the upper and lower bound I use $[l(t), h(t)]$ to denote t's report. When it is convenient to refer to these separately I write $l(t)$ or $h(t)$. A pure strategy of the receiver is a function $a : M \to A$. The expected payoff of the receiver given a, m and $\mu(t|m)$ is $\int u^R(a,t)d\mu(t|m)$. The equilibrium concept considered is Perfect Bayesian Equilibrium (referred to as PBE or simply as an "equilibrium" in the remainder of the paper), defined as follows:

Definition 4.1. *A PBE of* Γ *is a receiver strategy* $a^*(m)$*, a sender strategy* $m^*(t)$ *and for each* $m \in M$ *beliefs* $\mu^*(t|m)$*, such that*

1. *For all* $m \in M$*,* $a^*(m) \in \underset{a \in A}{\arg\max} \int u^R(a,t)d\mu^*(t|m)$.
2. *For all* $t \in T$*,* $m^*(t) \in \underset{m \in M(t)}{\arg\max} U(a^*(m), m)$.
3. *For any report* m *sent in equilibrium the beliefs of the receiver* $\mu^*(t|m)$ *are obtained by applying Bayes rule. For any report* $[l,h]$ *not sent in equilibrium* $\mu^*(t|[l,h])$ *have support* $[l,h]$.

This is a standard definition, which requires that the receiver chooses the strategy that maximizes her payoff given her beliefs, that the sender chooses the strategy that maximizes payoffs given the strategy of the receiver, and that the beliefs of the receiver are rational.

4.4 Costly Reporting and Equilibrium

4.4.1 Separating Equilibria

In this subsection the focus is on separating equilibria (SE) of Γ. A separating equilibrium is a PBE in which each sender type sends a unique report, i.e. for all $t, t' \in T$ it holds that $m^*(t) \neq m^*(t')$. Since the receiver's beliefs are rational, in a separating equilibrium she always knows exactly who sent the report and can choose the optimal action given each sender type. Hence, in a separating equilibrium all the private information of the sender is transmitted to the receiver, so communication is perfect.

Milgrom's (1981) unraveling result implies that in the absence of reporting costs, i.e. with $k = 0$, all equilibria of Γ are separating. The logic is that types prefer identifying themselves, which is always possible at zero cost, over pooling with lower types. It also holds that in any equilibrium $l^*(t) = t$ for all t, since otherwise there are types that would mimic the reports of higher types.

On the other hand, when k is positive, a sender strategy in which $l(t) = t$ for all $t \in T$ is not necessarily an equilibrium. The problem is that this may become too expensive for some types. For example, $t = 1$ would have to send the very precise report $[1, 1]$ and if k is sufficiently large this will be too expensive. However, full separation can be accomplished in another way. Reporting costs can be reduced by sending some $l(t) < t$. This solves the problem of too high reporting costs, but it creates another one. If $l(t) < t$, then types in $[l(t), t)$ may be tempted to mimic the report of t. In this subsection it will be shown that it is possible to find a sender strategy such that these temptations do not arise. The reporting costs can be used to avoid this. In equilibrium, these costs work as a signaling device and full separation is accomplished through a combination of costly signaling and disclosure of information.

Before proceeding to the main result a lemma is derived, which is helpful in characterizing separating equilibria. In order to formulate the lemma two different terms are introduced. First, let

$$t^*(t) := \max\{\arg\max_{t' \in [0,t]} U(a^R(t'), [t', 1])\}. \tag{4.1}$$

To understand this term, suppose we have a separating equilbrium in which $m(t) = [t, 1]$ for all $t \in T$. In this case, the receiver chooses $a^R(t)$ when she receives $m(t)$. Therefore, each t's payoff equals $U(a^R(t), [t, 1])$. $t^*(t)$ is the type below or equal to t that obtains the highest payoff in this separating equilibrium. If there is more than one type that obtains this payoff, then $t^*(t)$ is the highest of these types. For example, $t^*(1)$ is the type that would obtain the highest payoff by reporting all his good news, given that he is identified by the receiver. Note that if $k = 0$ then, given the monotonicity assumptions, $t^*(t) = t$ for all t. With $k > 0$ it is possible that $t^*(t) < t$, in particular if k is large. In this case, identification by reporting all the good news is expensive for high types. Let

$$\varepsilon(t) := \{l \in [0, t] : U(a^R(t), [l, 1]) = U(a^R(t^*(t)), [t^*(t), 1])\}. \tag{4.2}$$

As can be seen, $\varepsilon(t)$ is related to $t^*(t)$. Consider again a situation in which all sender types are identified by the receiver. Then $\varepsilon(t)$ is the lower bound of the report with upper bound 1 that t must send in order to obtain the same payoff as $t^*(t)$. In other words, $\varepsilon(t)$ is the amount of good news that t should report in order to obtain the same payoff as the type $t' < t$ that obtains the highest payoff by reporting all his good news. If $k = 0$, then $\varepsilon(t) = t$ for all t. With $k > 0$, it is possible that $\varepsilon(t) < t$. The following lemma shows that $\varepsilon(t)$ exists and is unique.

Lemma 4.1. *For all $t \in T$ it holds that $\varepsilon(t)$ exists and is unique. Further $t^*(t) \leq \varepsilon(t) \leq t$ and $\varepsilon(t)$ is strictly increasing in t.*

Proof. Suppose $t^*(t) = t$. Then, $\varepsilon(t) = t$ so it both exists and is unique.

Suppose $t^*(t) < t$, which means $U(a^R(t), [t, 1]) < U(a^R(t^*(t)), [t^*(t), 1])$. Then $u^S(a^R(t)) > u^S(a^R(t^*(t)))$ and $C(1-t) > C(1-t^*(t))$. This means that $U(a^R(t), [t^*(t), 1]) > U(a^R(t^*(t)), [t^*(t), 1])$. Then, by the continuity of C and the intermediate value theorem there exists some $\varepsilon(t) \in (t^*(t), t)$ such that $U(a^R(t), [\varepsilon(t), 1]) = U(a^R(t^*(t)), [t^*(t), 1])$.

Uniqueness follows from the fact that C is decreasing (and hence monotone).

Consider some $t' > t$. If $t^*(t') > t$ then $\varepsilon(t') > \varepsilon(t)$ since $\varepsilon(t) \leq t < t^*(t') \leq \varepsilon(t')$. If $t^*(t') \leq t$ then $t^*(t) = t^*(t')$. Then since $u^S(a^R(t)) < u^S(a^R(t'))$ we have $U(a^R(t^*(t')), [t^*(t'), 1]) = U(a^R(t), [\varepsilon(t), 1]) < U(a^R(t'), [\varepsilon(t), 1])$.This means that $\varepsilon(t') > \varepsilon(t)$ is needed for $U(a^R(t^*(t')), [t^*(t'), 1]) = U(a^R(t'), [\varepsilon(t'), 1])$. Hence, $\varepsilon(t)$ is increasing in t. □

Lemma 1 implies that if all types t send $[t, 1]$ and are identified by the receiver, then t' can always obtain the same payoff as the lower types with the highest payoff by sending $[\varepsilon(t'), 1]$. It is convenient to introduce the following notation, let $T_1 = \{t : \varepsilon(t) = t\}$and $T_2 = T \backslash T_1$. The following Proposition uses Lemma 1 to establish the existence of a separating equilibrium of Γ:

Proposition 4.1. *The following is a separating equilibrium of* Γ

1. *All* $t \in T$ *send* $m^*(t) = [\varepsilon(t), 1]$.

2. *For any* $[l, h]$ *not sent in equilibrium* $\mu(t|[l, h])$ *is degenerate at* l.

3. *For all* $t \in T$ *it holds that* $a^*(m^*(t)) = a^R(t)$, *and* $a^*([l, h]) = a^R(l)$ *for any message* $[l, h]$ *not sent in equilibrium.*

Proof. We first make three observations which help to prove the result:

(i) Lemma 1 guarantees that (1) defines a unique report for all $t \in T$.

(ii) Equilibrium payoffs are weakly increasing in t. This is because each type in equilibrium obtains $U(a^R(t^*(t)), [t^*(t), 1]) = \max\limits_{t' \in [0,t]} U(a^R(t'), [t', 1])$. Since t defines the upper bound of the domain over which this function is maximized, equilibrium payoffs are weakly increasing in t.

(iii) Claim: If $U(a^R(t), m^*(t)) > U(a^R(t'), m^*(t'))$ then $l(t) > t'$.

Proof: Assume $U(a^R(t), m^*(t)) > U(a^R(t'), m^*(t'))$. First, note that $U(a^R(t), m^*(t)) > U(a^R(t'), m^*(t'))$ implies $t > t'$ since equilibrium payoffs are increasing in t.

If $t \in T_1$ then $l(t) = t > t'$, which means that for $t \in T_1$ the claim holds.

If $t \in T_2$ then $l(t) = \varepsilon(t) > t^*(t) = \max\{[0,t] \cap T_1\}$. For all $t' \in [t^*(t), t]$ we have $U(a^R(t), m^*(t)) = U(a^R(t'), m^*(t'))$ since $[t^*(t), t] \subset T_2$. So for $[t^*(t), t]$ the hypothesis of the claim is not satisfied and we need not prove anything. On the other hand $l(t) = \varepsilon(t) > t^*(t)$ means that $l(t) > t'$ for all $t' \in [0, t^*(t)]$ and so the claim holds for any $t \in T_2$.□

Since equilibrium payoffs are increasing in t, no type t has incentives to deviate to a report $m^*(t')$ sent by some $t' < t$. At the same time (iii) implies that t can't deviate to any message $m^*(t')$ sent by some type $t' > t$, since $l(t') > t$ and therefore $m^*(t') \notin M(t)$.

It remains to show that no type has incentives to deviate to some report not sent in equilibrium. Each $t \in T$ obtains $U(a^R(t^*(t)), [t^*(t), 1])$ in equilibrium. Given the skeptical beliefs of R, a type t deviating to some out of equilibrium report $m^D = [l^D, h^D]$ will be identified as type $l^D < t$. The buyer will thus take action $a^R(l^D)$ in response to m^D and the payoff that t obtains from the deviation will be $U(a^R(l^D), [l^D, h^D]) \leq U(a^R(l^D), [l^D, 1]) \leq U(a^R(t^*(t)), [t^*(t), 1])$, where the second inequality follows from the definition of $t^*(t)$. Hence, there are no incentives to deviate to out of equilibrium reports. □

The fact that a separating equilibrium always exists means that unraveling always can occur even if reporting is costly. In fact, it can occur for arbitrarily large k, i.e. no matter how high the reporting costs are.

Since $l^*(t) = \varepsilon(t)$ for all t in the equilibrium specified in Proposition 4.1, the types in T_1 send $l^*(t) = t$ and the types in T_2 send $l^*(t) < t$. Hence, the types in T_1 behave as when reporting is costless and prove that they are not of a lower type than what they are. For these types, the benefit of separating is not outweighed by the cost of proving their type and they can therefore separate from lower types by disclosing information. However, for types in T_2 this strategy is too expensive and would give a lower payoff than some lower types that they could mimic. In equilibrium, they therefore economize on reporting costs by sending $l^*(t) = \varepsilon(t) < t$ and only prove that they are not lower than $\varepsilon(t) < t$. This means that enough costs

Figure 4.1: A Separating Equilibrium with Positive Reporting Costs

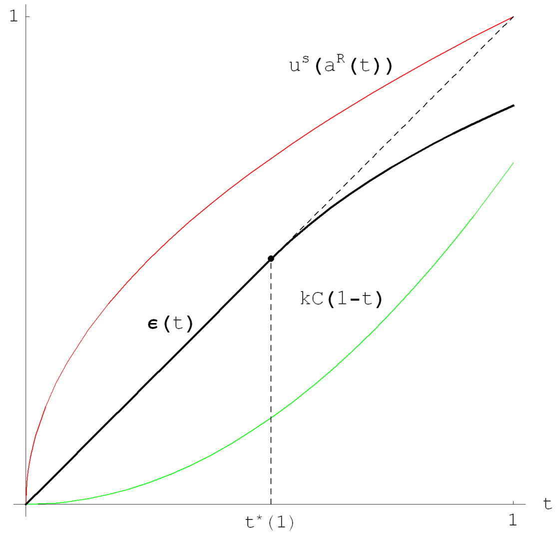

must be incurred so no type in $[\varepsilon(t), t]$ have incentives to mimic the report, which is accomplished by letting $l^*(t) = \varepsilon(t)$. In this way, types in T_2 obtain the same payoff as the type below them with the highest payoff. The costs incurred producing the report hence work as a signaling device. When it is too expensive to prove your type, it is instead possible to incur costs and in this way signal your type. Hence, in the equilibrium specified in Proposition 4.1, a combination of disclosure of information and costly signaling accomplishes full separation.

Figure 4.1 illustrates a separating equilibrium, in a setting in which $u^R(t,a) = t\log(a) - a$, $u^S(a) = a^{1/2}$ and $C(x) = 0.7(1-x)^2$. The $\varepsilon(t)$ curve gives the lower bound of the reports. Since $T_1 = [0, t^*(1)]$, $\varepsilon(t)$ is a 45 degree line in $[0, t^*(1)]$. Thereafter $\varepsilon(t)$ is strictly concave and therefore below the 45 degree line, which reflects the fact that types in $T_2 = [t^*(1), 1]$ must economize on reporting costs in equilibrium.

A couple of different properties of the equilibrium in Proposition 4.1 are worth pointing out. First, the sender's equilibrium payoffs are weakly increasing in t. The lowest payoff obtained in equilibrium is $U(a^R(0), [0,1])$ and the highest payoff is $U(a^R(t^*(1)), [t^*(1), 1])$. In segments of the type-space such that $l^*(t) = t$ payoffs are strictly increasing, while in segments of the type-space in which $l^*(t) < t$, payoffs are constant. This contrasts somewhat with the case in which reporting is costless, in which equilibrium payoffs are strictly increasing in t. Second, the precision of the reports is strictly increasing in t, meaning that higher types incur strictly higher costs writing the reports. This is simply because high types distinguish themselves from lower types by reporting a larger amount of good news. Thus, their reports contain more information and are therefore costlier to produce. Third, whereas the beliefs with respect to unsent reports are skeptical, the beliefs with respect to reports sent in equilibrium are not necessarily skeptical. More precisely, if T_2 is nonempty, the beliefs with respect to the messages sent by types in T_2 are non-skeptical. This contrasts with costless reporting, where beliefs are always skeptical. The intuition is that with costly reporting, not reporting all favorable information can be justified on the ground that it is too costly to do so, and hence the receiver can be somewhat less skeptical with respect to withheld information. Therefore, high reporting costs mute skepticism.

The equilibrium characterized in Proposition 4.1 is not the only separating equilibrium of Γ. For example, an equilibrium can be constructed in which types in T_2 continuously contract the reports at the upper end of the type space rather than at the lower end (as in Proposition 4.1). This raises the question if some separating equilibria economize more on reporting costs than others. For example, one may wonder whether there are equilibria in which types in T_1 choose some $l^*(t) < t$, or types in T_2 choose some $l^*(t) < \varepsilon(t)$, thereby saving on reporting costs. Or on the other hand, if an equilibrium of "dissipative signaling" in which the payoffs are constant over the entire typespace is possible. In the following proposition, it is shown that this is not the case. On the contrary, all sender types earn exactly the same payoff in any separating equilibrium. Formally, let two equilibria be *payoff equivalent* if all sender types and the receiver earn the same payoff in both equilibria. Then:

Proposition 4.2. *All separating equilibria of* Γ *are payoff equivalent.*

Proof. In the equilibrium in Proposition 4.1 all types $t \in T$ obtain $U(a^R(t^*(t)), [t^*(t), 1])$. We will show that this payoff must be obtained by all types in any SE.

Step (i): In any SE $U(a^R(t), [l(t), h(t)]) \geq U(a^R(t^*(t)), [t^*(t), 1])$ for all $t \in T$.

Proof: Consider a SE in which $U(a^R(t), [l(t), h(t)]) < U(a^R(t^*(t)), [t^*(t), 1])$ for some $t \in T$. If $[t^*(t), 1]$ is sent by some other type t' (which requires $t' \geq t^*(t)$) in equilibrium then t obtains $U(a^R(t'), [t^*(t), 1]) \geq U(a^R(t^*(t)), [t^*(t), 1])$, if he deviates to $[t^*(t), 1]$, so $[t^*(t), 1]$ is a profitable deviation.

If $[t^*(t), 1]$ is not sent in equilibrium then given the worst possible inference on account of R type t obtains a payoff $U(a^R(t^*(t)), [t^*(t), 1])$ if he deviates to $[t^*(t), 1]$. Hence $[t^*(t), 1]$ is a profitable deviation.

Step (ii): In any SE $U(a^R(t), [l(t), h(t)]) \leq U(a^R(t^*(t)), [t^*(t), 1])$ for all $t \in T$.

Proof: Consider a SE in which $U(a^R(t^0), [l(t^0), h(t^0)]) > U(a^R(t^*(t^0)), [t^*(t^0), 1])$ for some $t^0 \in T$. This implies that $l(t^0) < \varepsilon(t^0)$. In equilibrium it is then needed that $U(a^R(t), [l(t), h(t)]) \geq U(a^R(t^0), [l(t^0), h(t^0)])$ for all $t \in [l(t^0), t^0]$. However, this means that the same must apply for all types in $[l(l(t^0)), t^0]$ with $l(l(t^0)) < \varepsilon(l(l(t^0))) \leq \varepsilon(l(t^0)) \leq l(t^0)$. Iterating another step it must also hold for all types in $[l(l(l(t^0))), t^0]$. By iterating this argument it follows that a necessary condition for $m(t^0) = [l(t^0), h(t^0)]$ to be part of a SE is that there exists a sequence

$$\{(l(t^n), h(t^n)\}_{n=1}^{N},$$

with

$$\begin{aligned}
0 &\leq l(t^n) \leq t^n \leq h(t^n) \leq 1, \\
t^n &= l(t^{n-1}), \\
U(a^R(t^n), [l(t^n), h(t^n)]) &\geq U(a^R(t^0), [l(t^0), h(t^0)]) \ \forall n = 1, 2, ..., N, \\
\text{if } N \text{ is finite then } l(t^N) &= t^N
\end{aligned}$$

The last property follows since otherwise we would have to iterate the sequence an additional step. It is immediate to derive some additional properties that this sequence must satisfy:

$$\begin{aligned} l(t^n) &< \varepsilon(t^n), \\ l(t^n) &< l(t^{n-1}), \\ N &\to \infty. \end{aligned}$$

$l(t^n) < \varepsilon(t^n)$ follows since for any $h(t^n)$ we have $U(a^R(t^n), [\varepsilon(t^n), h(t^n)]) \leq U(a^R(t^*(t^0)), [t^*(t^0), 1]) < U(a^R(t^0), [l(t^0), h(t^0)])$. Next, as a consequence $l(t^n) < l(t^{n-1})$ since $l(t^n) < \varepsilon(t^n) \leq t^n = l(t^{n-1})$. Finally, since $l(t^n) < t^n$ for all $n = 1, 2, ..., N$, it must be that N is not finite. It then follows that $\{l(t^n)\}_{n=1}^{\infty}$ must converge to some limit $L \in [0, t^0]$. In what follows it will be shown that a sequence $\{l(t^n)\}_{n=1}^{\infty}$ satisfying the above properties cannot converge to some $L \in [0, t^0)$, which implies that there can be no separating equilibrium in which some type obtains a payoff higher than in the equilibrium in 4.1.

Claim: The sequence $\{l(t^n)\}_{n=1}^{\infty}$ does not converge to some $L \in [0, t^0)$.

Proof: Let $\Delta U(t^0) = U(a^R(t^0), [l(t^0), h(t^0)]) - U(a^R(t^*(t^0)), [t^*(t^0), 1]) = C(1 - \varepsilon(t^0)) - C(h(t^0) - l(t^0))$. That is, $\Delta U(t^0)$ is the increment in $t^{0\prime}s$ payoff over that obtained in the equilibrium in Proposition 4.1. Then for any t^n it must hold that $U(a^R(t^n), [l(t^n), h(t^n)]) - U(a^R(t^*(t^0)), [t^*(t^0), 1]) \geq \Delta U(t^0)$ and this implies $U(a^R(t^n), [l(t^n), h(t^n)]) - U(a^R(t^*(t^n)), [t^*(t^n), 1]) \geq \Delta U(t^0)$, or

$$C(1 - \varepsilon(t^n)) - C(h(t^n) - l(t^n)) \geq \Delta U(t^0) \ \forall t^n.$$

This expression defines an upper bound on $l(t^n)$ given $h(t^n)$ for all t^n which can be written $l(t^n) \leq h(t^n) - C^{-1}(C(1 - \varepsilon(t^n)) - \Delta U(t^0))$, where $C^{-1}(\cdot)$ is the inverse function of $C(\cdot)$, which is defined given that $C(\cdot)$ is decreasing. We can write

$$\begin{aligned} l(t^n) &\leq h(t^n) - C^{-1}(C(1 - \varepsilon(t^n)) - \Delta U(t^0)) \\ &\leq 1 - C^{-1}(C(1 - \varepsilon(t^n)) - \Delta U(t^0)) \\ &= \varepsilon(t^n) + [C^{-1}(C(1 - \varepsilon(t^n))) - C^{-1}(C(1 - \varepsilon(t^n)) - \Delta U(t^0))] \end{aligned}$$

Let $C^{-1}(C(1 - \varepsilon(t^n))) - C^{-1}(C(1 - \varepsilon(t^n)) - \Delta U(t^0)) := B(t^n, \Delta U(t^0))$ and note that $B(t^n, \Delta U(t^0)) < 0$. This means that $l(t^n)$ must be smaller than $\varepsilon(t^n)$ at least by the amount $B(t^n, \Delta U(t^0))$. Note that $B(t^n, \Delta U(t^0))$ cannot approach zero for some fixed $\Delta U(t^0)$. This is so since $\lim_{x \to y} C^{-1}(x) = \lim_{x \to y} C^{-1}(x - \varepsilon) \to C^{-1}(x) = C^{-1}(x - \varepsilon)$ by continuity and the latter equality does not hold since $C(\cdot)$ is decreasing. The consequence is that each t^n must send some $l(t^n)$ which fall short of t^n by some non-vanishing amount.

Now, suppose $\lim_{n \to \infty} l(t^n) = L$ for some $L \in [0, t^0]$. But the fact that $B(t^n, \Delta U(t^0))$ does not approach 0 implies that if we take some $t^{\hat{n}}$ sufficiently close to L we obtain that $l(t^{\hat{n}}) < L$. Hence, since $\{l(t^n)\}$ is a decreasing sequence $\lim_{n \to \infty} l(t^n)\} \neq L$ and the sequence can therefore not converge to some $L \in [0, t^0]$.□

We have therefore reached a contradiction, so in a separating equilibrium it cannot hold that $U(a^R(t^0), [l(t^0), h(t^0)]) > U(a^R(t^*(t^0)), [t^*(t^0), 1])$ for some $t^0 \in T$. □

The implication of this result is that any separating equilibrium will be similar to the one in Proposition 4.1. There can be some variation in the exact content of the reports, but the basic mechanism is the same. In equilibrium, types that cannot send $[t, 1]$ and be better off than lower types doing the same, must obtain payoffs equal to the highest one obtained by lower types. Types that are better off than all lower types by sending $[t, 1]$ must send precisely this message.

It is fairly straightforward to see that in a separating equilibrium no sender type can earn a lower payoff than the one obtained in Proposition 4.1. It is more difficult to see that no type can earn a higher payoff. The logic of the proof is the following. For a type to earn a higher payoff than in Proposition 4.1 he must produce a report with $l(t) < \varepsilon(t)$. But then, all types in $[l(t), \varepsilon(t)]$ can mimic this message and hence have to earn at least the same payoff. This means that the type $l(t)$ must produce a report with $l(l(t)) < \varepsilon(l(t))$. This creates a sequence that at some point hits zero, and for $t = 0$ it is not possible to produce a report with $l(0) < \varepsilon(0) = 0$. Therefore in a separating equilibrium it is not possible for any type to earn a higher payoff than in Proposition 4.1. The intuition then is that the precision of the message must be low enough to ensure at least as high payoff as that of lower types. But it must also be high enough to discourage lower types from mimicking it.

In the absence of reporting costs, i.e. when $k = 0$, there is a unique separating equilibrium (Milgrom 1981). The following result shows that a sufficient condition for all equilibria to be separating is that k is small enough.

Proposition 4.3. *For k sufficiently small all equilibria of Γ are separating.*

Proof. Consider an equilibrium in which some pooling occur. Suppose that in this equilibrium $\tau \subset T$ pool on some report $m^P = [l(\tau), h(\tau)]$, where evidently $l(\tau) \leq \min\{t \in \tau\}$. The payoff obtained for the pooling types in equilibrium is then $U(a^R(t(\tau)), [l(\tau), h(\tau)]$, where $a^R(t(\tau)) = \arg\max_{a \in A} \int u^R(a,t) d\mu(t|m^P)$. I write the optimal action of R as $a^R(t(\tau))$ in recognition of the fact that this action corresponds to the action R would take when certain that the message came from some type $t(\tau)$ in the convex hull of τ.

Consider a deviation to $[t, 1]$ by some $t \in \tau$ such that $t > t(\tau)$. The deviator t then earns at least $U(a^R(t), [t, 1])$. This deviation is then profitable if $u^S(a^R(t)) - u^S(a^R(t(\tau)) - k(C(1-t) - C(h(\tau) - l(\tau))) > 0$. The first term in this expression is positive since $t > t(\tau)$. If the second term is positive the deviation is therefore profitable. If the second term is negative, then for k sufficiently small $k(C(1-t) - C(h(\tau) - l(\tau))) < u^S(a^R(t)) - u^S(a^R(t(\tau))$. The deviation is therefore profitable for small enough k.

Hence, for k sufficiently small there can be no pooling in equilibrium, so all equilibria are separating. $\square$

Proposition 4.3 implies that for small values of k the unraveling result of Milgrom (1981) is intact, in the sense that all equilibria are separating. The unraveling result is thus fully robust to small reporting costs. This conclusion contrasts with that of Jovanovic (1981) and Verecchia (1983), in which full separation is impossible even with arbitrarily small reporting costs. As mentioned, the difference arises since in those papers the disclosure decision is binary, which implies that there are always some types for which disclosure is not worthwhile.

To sum up, in this subsection it has been established that separating equilibria exist, regardless of the reporting costs. All separating equilibria are payoff equivalent. If the reporting costs are sufficiently small, all equilibria are separating. In other words, the results in this subsection show how the unraveling result of Milgrom (1981) generalizes to reporting costs that increase continuously in the precision of the report.

4.4.2 Pooling Equilibrium

In a pooling equilibrium all types send the same report, which means that communication is absent. Hence, the receiver's information is not enhanced beyond her knowledge of the prior distribution and she chooses the same action regardless of the type of the sender. Here, in a pooling equilibrium, all types must send $[0,1]$, since this is the only report that is true for all sender types. Proposition 4.3 implies that when reporting costs are small, a pooling equilibrium does not exist. The following result shows that on the other hand, with high reporting costs, a pooling equilibrium exists.

Proposition 4.4. *For k sufficiently high a pooling equilibrium exists. In the pooling equilibrium $m^*(t) = [0,1]$ for all $t \in T$.*

Proof. (i) In a pooling equilibrium all types must send $[0,1]$ since this is the only report available to all types. Hence, if a pooling equilibrium exists it is unique.

(ii) Assume all types send $m^P = [0,1]$. This means that in equilibrium their payoff is
$U(a(m^P),[0,1])$, where $a(m^P) = \arg\max_{a\in A} \int u^R(a,t)dF(t)$. Note that $a(m^P) = a^R(\hat{t})$ for some $\hat{t} \in (0,1)$. If the beliefs of R with respect to out of equilibrium reports are skeptical, then no type in $[0,\hat{t}]$ can profitably deviate, since the action of the receiver in response to the deviation would be lower than $a(m^P)$ and the deviating report would be costlier than $[0,1]$.

Profitable deviations are available for types $t > \hat{t}$ such that $U(a^R(\hat{t}),[0,1]) < U(a^R(t),[t,1])$ or $u^S(a^R(t)) - u^S(a^R(\hat{t})) > k(C(1-t) - C(1))$. The left side of this inequality is bounded and $C(1-t) - C(1)$ is bounded away from zero since $t > \hat{t}$. This means that for k sufficiently large the inequality will not hold for any t, so there will be no profitable deviation. Therefore, for sufficiently high k all types sending m^P together with a skeptical inference on account of the receiver is an equilibrium. □

A pooling equilibrium arises when the reporting costs are sufficiently high, since a deviation in which the sender reports enough information to prove his type becomes too expensive. Skeptical beliefs with respect to unsent reports add to these costs.[8] One may wonder why separating equilibria exist regardless of the reporting costs, but still it becomes impossible to break out of a pooling equilibrium when the costs are high. The reason is that in the separating equilibrium, non-skeptical equilibrium beliefs of the receiver allow the sender to identify himself by providing less informative reports when the costs are high. In contrast, in a pooling equilibrium the receiver may be skeptical with respect to all unsent reports. In this case, the sender would have to provide enough evidence to overcome the receivers skepticism, i.e., to identify himself t must send $[t, 1]$. Hence, in breaking out from an equilibrium of uninformative reports, the sender cannot justify withheld information by appealing to the high costs of providing it, as in the separating equilibrium. This makes it much harder for him to identify himself.

The fact that k high enough is sufficient for a pooling equilibrium to exist, means that k small enough is not only sufficient for separating equilibria to be the unique class of equilibria, it is also necessary. In principle, equilibria that constitute intermediate cases between pooling and full separation may also arise, such as equilibria with multiple pools or equilibria in which some types pool and others separate. Here, I have limited myself to show that uninformative equilibria arise when the reporting costs are sufficiently high. Future work could analyze intermediate cases more closely.

In the pooling equilibrium, the sender incurs no cost writing reports. It can be shown that whenever a pooling equilibrium exists, the sender prefers the pooling

[8] Recall that a belief $\mu(t|[l, h])$ is skeptical if it is degenerate at l.

equilibrium. In fact, all sender types prefer the pooling equilibrium to the separating equilibrium. This is illustrated by the following argument. The maximum payoff earned in a separating equilibrium is $U(a^R(t^*(1)), [t^*(1), 1])$. At the same time, given the skeptical beliefs of the receiver, $t^*(1)$ could obtain at most $U(a^R(t^*(1)), [t^*(1), 1])$ by deviating from the pooling equilibrium to $[t^*(1), 1]$. But since we are in an equilibrium, this is not a profitable deviation, so $U(a^R(t^*(1)), [t^*(1), 1]) < U(a^R(\hat{t}), [0, 1])$.[9] Hence, the payoffs obtained in the pooling equilibrium exceed the maximum (over all types) payoff obtained in the separating equilibrium. Therefore all sender types prefer the pooling equilibrium and consequently so does the sender.

On the other hand, the receiver always prefers the separating equilibrium. It is therefore not immediately clear whether the pooling or the separating equilibrium is better in welfare terms. Such a comparison depends on a cardinal comparison of $u^S(a)$ and $u^R(a, t)$. In fact, not even when the intensity k of the reporting cost is arbitrarily high is it possible to say which of the equilibria is better in welfare terms. The reason is that the payoff the sender obtains in a separating equilibrium is bounded from below at $U(a^R(0), [0, 1])$. Hence, the sender's gain by going from a separating equilibrium to a pooling equilibrium is bounded at $U(a^R(\hat{t}), [0, 1]) - U(a^R(0), [0, 1])$.

4.5 A more Active Receiver

In this section an extension of the baseline model of costly reporting presented in Section 3 is considered. In this extension, the receiver must examine the report at a cost in order to access the information contained in it. The aim is to capture the fact that just as it is costly to elaborate a report, it frequently requires both time and effort to read and understand it.

When the receiver actively decides whether to read a report, it becomes important whether she can make this decision contingent on some first impression of the report's appearance. This first impression may come from a first glance or a quick

[9] With $\hat{t}$ defined as in the proof of Proposition 4.4.

browse. It is then relevant how the first impression is related to the information content of the report and to which extent it can be manipulated. I will assume that the appearance of a report is related to its precision. The idea is that since a precise report contains more information it tends to be thicker and denser. It may also require more numbers, arguments, explanations, diagrams and other details that can be detected by a quick browse. I also assume that appearances can be manipulated at a cost. Hence, a report can be "polished" to look more or less precise than what it really is. This captures the idea that, for example, a vague report can be made to appear more informative at a first glance by filling it with nonsense information (that nevertheless looks informative at first glance). Conversely, a precise report can be condensed so it seems less informative, perhaps by making an additional effort to eliminate redundant information.

4.5.1 The Model

The same model as the one in Section 3 is considered, but with some modifications. Each sender type t chooses an element $m \in M(t)$ as before, but now he also chooses an *appearance* $p \in P = [0, 1]$. Hence, here a sender of type t chooses a couple $(m, p) \in M(t) \times P$. In this section I will refer to the couple (m, p) as a report, to m as the *information content* of (m, p) and to p as the *appearance* of (m, p). If p is chosen so $p < \upsilon(m)$ the report is made to look more precise than what it is and if $p > \upsilon(m)$, it is made to look less precise than what it is. If $p = \upsilon(m)$, the appearance is not manipulated, so the report looks just as precise as it really is. The interpretation is the following. A precise report here, is understood as a report that contains a large amount of information. Reports of a given precision may have a things in common. For example, they may be equally thick, or contain similar amounts of numbers, plots, calculations and arguments, and so on. The sender can make a report look more precise than what it is by including redundant information. He could also make it look more precise than what it is, by being more concise and eliminating redundancies.

The cost of producing a report (m,p) is given by $kC(\upsilon(m))+k^p\psi(\upsilon(m)-p)$, where the first term is the cost function from Section 3, ψ is a function $\psi : [-1,1] \to \mathbb{R}$ and $k^p > 0$ is an intensity parameter. The term $k^p\psi(\upsilon(m)-p)$ gives the cost of manipulating the appearance of a report with information content m so it looks like a report of precision p. The function ψ is assumed to be a convex and satisfy $\psi'(x) < 0$ for $x < 0$, $\psi'(x) > 0$ for $x > 0$ and $\psi(0) = 0$[10]. This means that it is costly to make a report look both more or less precise than what it is, i.e. it is costly both to add redundant information and remove it. If the report is not manipulated, no additional costs are incurred. Additionally, in this section it will be assumed that C is differentiable and convex.

When a report (m,p) is delivered to the receiver, she immediately observes p. She then forms beliefs $\mu_1((m,t)|p)$. These beliefs are a probability distribution over $M \times T$. In other words, the receiver forms a belief both with respect to which types might have sent a report with appearance p, and with respect to the information content of the report. Next, she chooses an effort level $e \in E = [0,1]$ with which she examines the report. With probability e she understands the report and accesses the information contained in it, i.e. she observes m. With probability $1-e$ she does not access the information. If she accesses m, she updates her beliefs to $\mu_2(t|(m,p))$ and chooses an action $a \in A$. If she does not access the information, she does not update her beliefs, but just chooses an action $a \in A$. Examination effort costs $k^R\gamma(e)$, where γ is an increasing and convex function $\gamma : E \to \mathbb{R}$ and k^R parameterizes the intensity of the examination cost.

The sender's payoff is now $U(a,m,p) = u^S(a) - kC(\upsilon(m)) - k^p\psi(\upsilon(m)-p)$ and the receiver's payoff is $u^R(a,t) - k^R\gamma(e)$. The same assumptions as in Section 3 apply to these functions.

The sender's strategy is a function $\sigma^S : T \to M \times P$ with the restriction $m \in M(t)$. Then, $\sigma^S(t) = (m,p)$ means the sender of type t chooses report m and appearance p. I let $m(t)$ and $p(t)$ respectively denote the information and appearance components of $\sigma^S(t)$.

[10] However, it is not assumed that $\psi\prime(0) = 0$.

The receiver's strategy can be written as a combination of two functions. Let $\sigma_1^R : P \to E \times A$ and $\sigma_2^R : M \times P \to A$, where σ_1^R assigns an effort level and an action to each $p \in P$ and σ_2^R assigns an action to each $(m,p) \in M \times P$. Then $\sigma_1^R(p) = (e,a)$ means that the receiver chooses effort level e when receiving a report with appearance p, and action a if the information content is not accessed. Similarly $\sigma_2^R(m,p) = a$ means that the receiver chooses a in response to a report (m,p) (given that m is accessed). I write $e(p)$ and $a(p)$ to denote the effort and action components of σ_1^R. The strategy of the receiver is written as a pair (σ_1^R, σ_2^R).

A PBE here is a receiver strategy $(\hat{\sigma}_1^R, \hat{\sigma}_2^R)$, a sender strategy $\hat{\sigma}^S$ and beliefs $\hat{\mu}_1((m,t)|p)$ and $\hat{\mu}_2(t|(m,p))$, such that (i) $\hat{\sigma}_2^R$ maximizes the receiver's payoff given $\hat{\mu}_2(t|(m,p))$, (ii) $\hat{\sigma}_1^R$ maximizes the receiver's payoffs given $\hat{\mu}_1((m,t)|p)$, (iii) $\hat{\sigma}^S$ maximizes the sender's payoff given $(\hat{\sigma}_1^R, \hat{\sigma}_2^R)$ and (iv) $\hat{\mu}_1((m,t)|p)$ and $\hat{\mu}_2(t|(m,p))$ are rational (i.e. consistent with the receiver's strategy and the prior distribution).

4.5.2 Results

The model just described has multiple PBE. Here, I focus on separating equilibria. The aim is to study to what extent communication is complicated by the fact that the sender has to make an effort in order to assimilate the information contained in a report. One way of doing this is by focusing on the possibilities for perfect communication under these circumstances. Two categories of equilibria are identified and characterized: *Non-Reading Separating Equilibria* (NRSE) and *Reading Separating Equilibria* (RSE). A non-reading separating equilibrium is a separating PBE in which the receiver's examines any report with zero effort. A reading separating equilibrium is a separating PBE in which the receiver examines some messages sent in equilibrium with positive effort.

Non-Reading Separating Equilibrium

In a NRSE the reader never exerts any effort in order to access the information content of the report. In spite of this, the equilibrium is fully separating, so she is able to distinguish between all sender types. For this to be possible, all reports must have a unique appearance, so the receiver can tell them apart by just a quick glance. The fact that the receiver does not read the reports means that the information content becomes less relevant. The receiver never actually assimilates the information content of a report. The relevance of the information content instead comes from the costs involved in providing it. In order to separate when the receiver does not read the reports, the different sender types must "burn money" producing the reports to the degree that all sender types obtain the same payoff. The reason is that since the receiver does not read the reports, it is always possible for any type to mimic the report of all other types. The necessary costs can be incurred both by providing verifiable information and polishing the appearance of the report. Since separation is accomplished through the costs of producing reports and not by disclosure of information, the NRSE is a pure costly signaling equilibrium.

Two classes of NRSE are identified, *truth-telling* and *inflated-talk* equilibria. In a truth-telling equilibrium the sender never makes an effort to manipulate the appearance of the report, so $\upsilon(m(t)) = p(t)$ for all t. In an inflated-talk equilibrium some sender types manipulate the appearance of the report in order to make it look more precise than what it is, so there are t such that $\upsilon(m(t)) > p(t)$.[11] Loosely speaking, truth-telling equilibria arise if it is relatively more expensive to polish the report's appearance so it looks more precise, than to actually increase its precision. Inflated talk equilibria arise in the opposite case, when it is relatively more costly to provide information than to polish the appearance of the report. On the other hand, in a NRSE no type ever incur costs to make the reports look less precise.

In a NRSE it must hold that $\sigma^S(0) = ([0,1],1]$, i.e. the lowest type must send the cheapest possible report. If he sends any other report, he could always deviate

[11] Inflated-talk equilibria bear resemblance to those characterized in Kartik (2009) and Kartik et. al. (2006). In these models the sender incurs lying costs when claiming to be a higher type than what he is. In equilibrium, all sender types claim to be of a higher type than what they are. A difference in the present paper is that the reference point (the precision of a report) with respect to which a statement (an appearance) can be considered to be a lie, is endogenous. I.e., whereas in Kartik (2009) and Kartik et. al. (2006), the term $\upsilon(m(t))$ is a message with exogenous meaning "my type is t", here it is endogenously chosen by the sender.

to the cheaper report $([0,1],1)$ and be identified at least as $t=0$. Since all types must obtain the same payoff in a NRSE, this implies that they must all obtain $U(a^R(0),[0,1],1)$. Let $U(0):=U(a^R(0),[0,1],1)$. I now define a function that will be helpful in characterizing NRSE. Let $p_t(x)$ be the implicit function defined by

$$\{p : kC(x)+k^p\psi(x-p)=u(a^R(t))-U(0) \text{ and } p\leq x\} \tag{4.3}$$

This means that if $\upsilon(m(t))=x$ then $p_t(x)$ gives the appearance that the report of type t must have in a separating equilibrium for his payoff to equal that of type 0. Therefore, $(m,p_t(\upsilon(m(t)))$ gives the pairs (m,p) that t could send in a NRSE. The function $p_t(x)$ plays an important role in establishing the following lemma that characterizes the set of NRSE.

Lemma 4.2. *1. A unique set of payoff equivalent NRSE exists iff* $k>\frac{u(a^R(1))-U(0)}{C(0)}$, $k^p>\frac{u(a^R(1))-U(0)}{\psi(1)}$ *and the domain of* $p_1(x)$ *is an interval* $[\underline{m}_1,1]$.

2. In the NRSE all $t\in T$ *send* $(m^*(t),p_t(\upsilon(m^*(t))))$, *with* $m^*(t)$ *such that* $\upsilon(m^*(t))=\underset{x\in[\underline{m}_1,1]}{\arg\min}\,p_t(x)$.

3. If $\upsilon(m^*(1))=\underline{m}_1$ *the NRSE are truth telling equilibria. If* $\upsilon(m^*(1))>\underline{m}_1$ *there is a set of types* $(\hat{t},1]$ *such that* $p_t(\upsilon(m^*(t)))<\upsilon(m^*(t))$ *for all* $t\in(\hat{t},1]$, *while* $p_t(\upsilon(m^*(t)))<\upsilon(m^*(t))$ *for all* $t\in[0,\hat{t}]$. *Hence, these NRSE are inflated talk equilibria.*

Proof. First note that there is no NRSE in which $\upsilon(m(t))<p(t)$ for some t. In this case, the payoff of t is $u(a^R(t))-k^p\psi(\upsilon(m(t))-p(t))-kC(\upsilon(m(t)))$ and a profitable deviation is $(m\prime,p(t))$ such that $\upsilon(m')=p(t)$. Hence, in any equilibrium $\upsilon(m(t))\geq p(t)$ for all $t\in T$.

Sufficiency of 1: Note that if the domain of $p_1(m)$ is an interval $[\underline{m}_1,1]$ then the domain of $p_t(m)$ is an interval $[\underline{m}_t,1]$ for all $t\in T$.

Suppose all $t\in T$ send $(m(t),p_t(\upsilon(m(t))))$ such that $\upsilon(m(t))\in\underset{x\in[\underline{m}_t,1]}{\arg\min}\,p_t(x)$. Then all types send different reports of different appearances and all types earn the same payoff $u(a^R(t))-k^p\psi(\upsilon(m(t))-p(\upsilon(m(t)))-kC(\upsilon(m(t)))$. No deviation to messages $m'\in[\underline{m}_t,1]$ is profitable, since it renders payoffs $u(a^R(t))-k^p\psi(\upsilon(m')-$

$p_t(\upsilon(m(t)))) - kC(\upsilon(m')) \leq u(a^R(t)) - k^p\psi(\upsilon(m') - p_t(\upsilon(m'))) - kC(\upsilon(m'))$

$= u(a^R(t)) - k^p\psi(\upsilon(m(t)) - p_t(\upsilon(m(t)))) - kC(\upsilon(m(t)))$, where the first inequality follows since $p(\upsilon(m(t))) = \min\limits_{x\in[\underline{m}_t,1]} p_t(x) \leq p(\upsilon(m'))$ and $\upsilon(m') \geq p_t(\upsilon(m(t)))$.

Deviations to messages in $m' \in [0, \underline{m}_t)$ are likewise unprofitable. First, note that $p_t(\underline{m}_t) = \underline{m}_t$ (since $C(0) > u(a^R(1)) - U(0)$). Next, for any deviation to $m' \in [0, \underline{m}_t)$ the profits obtained are $u(a^R(t)) - k^p\psi(\upsilon(m') - p_t(\upsilon(m(t)))) - kC(\upsilon(m')) \leq$

$u(a^R(t)) - k^p\psi(0) - kC(\upsilon(\underline{m}_t)) = u(a^R(t)) - k^p\psi(\upsilon(m(t)) - p_t(\upsilon(m(t)))) - kC(\upsilon(m(t)))$, where the first inequality follows since $\upsilon(m') - p_t(m(t)) \geq 0$ and $\upsilon(m') < \underline{m}_t$.

Hence, the assumptions in (1) guarantee that there is a NRSE in which all types send $(m(t), p(\upsilon(m(t))))$ such that $\upsilon(m(t)) = \arg\min\limits_{x\in[\underline{m}_t,1]} p_t(x)$.

Necessity of 1: We examine each of the assumptions in (1). First, note that in equilibrium all types must send $(m(t), p_t(\upsilon(m(t))))$ such that $\upsilon(m(t)) \in \arg\min\limits_{x\in D(p_t)} p_t(x)$, where $D(p_t)$ is the domain of p_t. If this is not the case, a deviation to m' such that $\upsilon(m') \in \arg\min\limits_{x\in D(p_t)} p_t(x)$ would be profitable.

Suppose $kC(0) < u(a^R(1)) - U(0)$. Then $\min\limits_{x\in D(p_t)} p_t(x) = 0$. Hence in equilibrium $t = 1$ must send $(m(t), 0)$. Now consider a deviation to $(m', 0)$ with $\upsilon(m') = 0$. This gives payoffs $u(a^R(1)) - C(0) < U(0) = u(a^R(1)) - k^p\psi(\upsilon(m(t))) - kC(\upsilon(m(t))))$. So profitable deviations are available.

Now suppose $k^p\psi(1) < u(a^R(1)) - U(0)$. Again in equilibrium $t = 1$ must send $(m(t), 0)$. Now consider a deviation to $(m', 0)$ with $\upsilon(m') = 1$. This gives payoffs $u(a^R(1)) - k^p\psi(1) < U(0) = u(a^R(1)) - k^p\psi(\upsilon(m(t))) - kC(\upsilon(m(t)))$. So profitable deviations are available.

Finally, suppose that the domain of $p_1(x)$ is not an interval $[\underline{m}_1, 1]$. This means that there is some point a in $[0,1]$ such that for some $\varepsilon \in (0, 1-a]$ it holds that $a+\varepsilon \notin D(p_1)$. Further, $p(a) = 0$. Consider a deviation to m' such that $\upsilon(m') = a+\varepsilon$. At this point $kC(a+\varepsilon) + k^p\psi(a+\varepsilon) < u(a^R(t)) - U(0)$ since if the inequality were reversed there would be some p at which it held with equality, contradicting that $a+\varepsilon \notin D(p_1)$. Hence $m' = a+\varepsilon$ is a profitable deviation.

2. It was already shown in the first paragraph of the proof of necessity of 1 that all types must send $m(t)$ such that $\upsilon(m(t)) \in \underset{x\in D(p_t)}{\arg\min} p_t(x)$ and $p(\upsilon(m(t)))$.

3. First, suppose $\underset{x\in[\underline{m}_1,1]}{\arg\min} p_1(x) = \underline{m}_1$. This evidently means that $t=1$ tells the truth in equilibrium. However, it also means that for $t \in (0,1)$ it holds that $\underset{m\in[\underline{m}_1,1]}{\arg\min} p_t(m) = \upsilon(m^*(t))$. To see this, suppose that for some $t \in (0,1)$ $\underset{m\in[\underline{m}_t,1]}{\arg\min} p_t(m) < \upsilon(m^*(t))$. In this case, due to the convexity of $p_t(m)$, it would have to hold that $p_t'(\underline{m}_t) < 0 < p_1'(\underline{m}_1)$. But it can be shown that $p_t'(\underline{m}_t) = \frac{k^p\psi'(0)+kC'(\underline{m}_t)}{k^p\psi'(0)}$ and $p_t'(\underline{m}_t) < p_1'(\underline{m}_1)$ is therefore equivalent to $C'(\underline{m}_1) < C'(\underline{m}_t)$ which does not hold since $\underline{m}_1 < \underline{m}_t$. Hence, if $t=1$ speaks the truth in equilibrium then so does all types (alternatively, $p_t'(\underline{m}_t)$ is increasing in t. Since $p_t(\underline{m}_t)$ is convex a positive slope means that all types will minimize $p_t(m)$ at $\underline{m}_t$.)

Now, suppose $\underset{x\in[\underline{m}_1,1]}{\arg\min} p_1(x) > \underline{m}_1$. Then $t=1$ will necessarily inflate talk in equilibrium. Further, there will be a set of types $(\hat{t},1]$ such that $\underset{x\in[\underline{m}_t,1]}{\arg\min} p_t(x) = \upsilon(m^*(t))$ for all $t \in (\hat{t},1]$. However, as argued, if for some $\hat{t}$ $\underset{x\in[\underline{m}_{\hat{t}},1]}{\arg\min} p_{\hat{t}}(x) = \underline{m}_{\hat{t}}$ then all types $[0,\hat{t}]$ will tell the truth. Indeed, there is always such a $\hat{t}$, namely 0, who always tells the truth, so the question is rather if some other types besides him tell the truth. The set of truth tellers will be the set of types t such that $p_t'(\underline{m}_t) > 0$. □

Lemma 4.2 characterizes the set of NRSE. Item (1) gives necessary and sufficient condition for the existence of a unique set of payoff equivalent NRSE. First, both k and k^p must be sufficiently high. The reason is that all types must be able to burn enough money providing information and manipulating appearances. However, this is not enough, the costs $kC(0)$ and $k^p\psi(1)$ must be high enough by themselves and the domain of $p_1(x)$ must be an interval. These requirements arise since for any report (m,p) sent in equilibrium, there can be no report (m',p), i.e. with the same appearance but different information content, that is less costly to provide than (m, p). Hence, given the equilibrium message, it must be more costly both to increase and decrease the (true) precision of the report. A point where this holds for all types exists under the conditions stated in (1) of Lemma 4.2. Item (2) states that (m, p) must be chosen so $\upsilon(m)$ minimizes $p_t(\upsilon(m)$ and $p = p_t(\upsilon(m))$. The reason is that such points are the only ones in which there is no report (m', p) with lower cost. The intuition is that at these points, the cost of increasing precision outweighs the lower manipulation costs, and the savings obtained by decreasing precision, are outweighed by higher manipulation costs. Item (2) also shows in which sense there is multiplicity of NRSE. The report of each type must have a unique precision, but there is an infinite number of reports with the same precision. Hence, the report of each type has a unique precision and appearance, but the actual information content is not determined. Finally, (3) gives conditions under which the NRSE are truth-telling or inflated talk equilibria.

Lemma 4.2 gives a set of conditions that are related to NRSE. However, these are not in terms of the primitives of the model. Whether the domain of p_1 is an interval $[\underline{m}_1, 1]$ depends on the properties of $k^p\psi(\cdot)$ and $kC(\cdot)$ and it has not been established so far whether this is at all possible, or whether it is consistent with the conditions imposed on $k^p$and k. However, it is helpful in establishing the main result of this subsection, in which the equilibria is characterized in terms of the intensity parameters of the cost functions.

Proposition 4.5. *If* $k > \frac{u(a^R(1))-U(0)}{C(0)}$ *then*

(1) If $k^p > -\frac{kC'(\underline{m}_t)}{\psi'(0)}$ *a unique set of payoff equivalent truth-telling NRSE exists.*

(2) There is some $\underline{k}^p(k)$ *such that for* $k^p \in [\underline{k}^p(k), -\frac{kC'(\underline{m}_t)}{\psi'(0)})$ *a unique set of payoff equivalent inflated talk NRSE exists, in which a set of types* $(\hat{t}(k^p), 1]$ *manipulate appearances and* $[0, \hat{t}(k^p))$ *does not.*

(3) For k^p *low enough there are no NRSE.*

Proof. (1) We can use lemma 4.2 and then only need to show that $\arg\min_{x\in[\underline{m}_1,1]} p_1(x) = \underline{m}_1$ if $k^p > \bar{k}^p(k)$ for some $\bar{k}^p(k) > 0$. Since $p_t(x)$ is convex, $\arg\min_{x\in[\underline{m}_1,1]} p_1(x) = \underline{m}_1$ if and only if $p_1'(\underline{m}_1) > 0$, i.e. if $k^p\psi'(0) + kC'(\underline{m}_t) > 0$ or $k^p > -\frac{kC'(\underline{m}_t)}{\psi'(0)}$.

(2) For an equilibrium to display inflated talk it is required that $\arg\min_{x\in[\underline{m}_1,1]} p_1(x) > \underline{m}_1$. A necessary condition for this is that $p_1'(\underline{m}_1) < 0$ or $k^p < -\frac{kc'(\underline{m}_t)}{\psi'(0)}$. It is also required that the domain of $p_1(x)$ be an interval $[\underline{m}_1, 1]$ i.e. there is no discontinuity in $[\underline{m}_1, 1]$. A sufficient condition for this is that $k^p\psi'(\underline{m}_1) > -kC'(\underline{m}_1)$ or $k^p > -k\frac{C'(\underline{m}_1)}{\psi'(\underline{m}_1)}$. To see that it is sufficient note that $kC(\upsilon(m)) + k^p\psi(\upsilon(m)) \geq u(a^R(1)) - U(0)$ for all m such that $\upsilon(m) \geq \underline{m}_1$ is sufficient for $p_1(x)$ to be defined on $[\underline{m}_1, 1]$. Next $kC(\underline{m}_1) + k^p\psi(\underline{m}_1) > u(a^R(1)) - U(0)$ and $k^p\psi'(\underline{m}_1) > -kC'(\underline{m}_1)$ implies $k^p\psi'(\upsilon(m)) > -kC'(\upsilon(m))$ for all m such that $\upsilon(m) > \underline{m}_1$ by convexity of $\psi(\cdot)$ and $C(\cdot)$. Hence $kC(\upsilon(m)) + k^p\psi(\upsilon(m)) \geq u(a^R(1)) - U(0)$ for all m such that $\upsilon(m) \geq \underline{m}_1$. Hence if $k^p \in [-k\frac{C'(\underline{m}_1)}{\psi'(\underline{m}_1)}, -k\frac{C'(\underline{m}_t)}{\psi'(0)}]$ there is an equilibrium that displays inflated talk. (Note that $k^p > -k\frac{C'(\underline{m}_1)}{\psi'(\underline{m}_1)}$ together with $k > \frac{u(a^R(1))-U(0)}{C(0)}$ implies the necessary condition for k^p, i.e. $k^p > \frac{u(a^R(1))-U(0)}{\psi(1)}$)

As is known by lemma 4.2 a set of types $(\hat{t}(k^p), 1]$ display inflated talk. This set can be more exactly specified by noting that $p_t'(\underline{m}_t) = p_t'(C^{-1}(U(t)/k))$ and thus $\frac{dp_t'(C^{-1}(U(t)/k))}{dt} = p_t''(C^{-1}(U(t)/k))\frac{1}{C'(t)}\frac{U'(t)}{k} < 0$. Hence even if $p_1'(\underline{m}_1) < 0$ we have that $p_t'(\underline{m}_t)$ will be on the increase in t so at some point $\hat{t}(k^p)$ it may become positive and then types $t > \hat{t}(k^p)$ will tell the truth. If $\lim_{t\to 0} p_t'(\underline{m}_t) = \frac{k^p\psi'(0)+kc'(1)}{k^p\psi'(0)} < 0$ which is equivalent to $k^p < -\frac{kc'(1)}{\psi'(0)}$ then all types except $t = 0$ inflate talk.

(3) This follows from lemma 4.2. □

Proposition 4.5 shows that k and k^p large enough is indeed a sufficient condition for NRSE to exist is. Moreover, if the manipulation cost is high enough relative to the cost of precision, all NRSE are truth-telling. If the manipulation cost is somewhat lower, there is a unique class of NRSE that display inflated talk. If the manipulation cost is too low, there are no NRSE (this is already known from Lemma 4.2 but is stated in Proposition 4.5 for completeness). The intuition behind this result is the following. In a NRSE it should not be possible for a type to maintain the same appearance but deviate to a report of different precision. If the cost of manipulation is high enough, relative to the precision costs, then truth telling can be an equilibrium since any deviation from such an equilibrium, maintaining appearance constant, will increase *manipulation* of appearance and therefore be too expensive. If the cost of manipulation is somewhat lower, there may be incentives to incur some manipulation cost in order to reduce the precision costs, which means truth telling is not an equilibrium. If the cost of manipulation is not too low, then due to the convexity of the manipulation and precision costs, it is possible to find a point from which the cost of additional manipulation outweighs the reduction of precision costs. Such a point is then part of an inflated talk equilibrium. If the cost of manipulation is too low, this point is never reached and it is always warranted to incur some additional manipulation costs in order to reduce the precision costs. In this case, there are no NRSE.

In a NRSE, higher types tend to send more precise reports, but also to incur higher costs polishing appearances. In fact, in an inflated talk NRSE, high types polish the reports to make them look more precise, whereas low types do not. The reason is that since low types send less precise messages, the convexity of the precision costs eventually leads the manipulation costs to become high relative to the precision costs for lower types.

The result in this subsection shows that communication can be possible even if the receiver never actually reads the report. In this case, communication is in terms of the appearances of the reports and the receiver can distinguish between the different reports simply by observing that they look different at first glance. For example, some reports may be thicker or appear to have more numbers, equations and so on. However, this communication comes at a high cost. All the benefits of being of a higher type is burnt up in communication costs. Nevertheless, while the sender should be unhappy with this state of affairs, it is an ideal situation for the receiver, who obtains full information without reading a single report.

It should be noted that there is always a non-reading pooling equilibrium in which all types send $([0, 1], 1)$ (a non-sense non-manipulated message) and the receiver never reads anything. This is worse for the receiver, but much better for all sender types.[12]

In the following subsection, the attention is turned to equilibria in which the receiver do read the reports.

Reading Separating Equilibrium

Under some circumstances there are separating equilibria in which reports are read. A few observations can be made about such equilibria. First, for the reader to examine a report with positive effort, there must be more than one type producing a report with the same appearance. Otherwise, in equilibrium the receiver knows already exactly what type sent the report and there are therefore no incentives to make an effort to understand it. Second, for a reading equilibrium to be separating any message that is examined with positive effort must be examined with effort $e = 1$, so the probability of interpreting the report correctly is equal to one. If this is not the case, the receiver sometimes is uncertain with respect to the sender's type and in these cases her choice will not correspond to that of full information. This means that full separation requires that the intensity of the reading costs is not too high. More precisely, when choosing effort level given a set τ of types pooling on some p in a separating equilibrium the receiver solves

$$\max_{e\in[0,1]}\{e\int_\tau u(a^R(t),t)f(t|\tau)dt + (1-e)\max_a \int_\tau u^R(a,t)f(t|\tau)dt - k^R\gamma(e)\}. \quad (4.4)$$

If τ consists of a single type then the solution to this problem is $e^* = 0$, i.e. the message is not read. If τ does not consist of a single type, the first order condition is $\int_\tau u(a^R(t))f(t|\tau)dt - \max\limits_a \int_\tau u(a^R(t))f(t|\tau)dt - k^R\gamma'(e^*) \geq 0$ if $e^* = 1$. Consequently, in a reading SE given any set τ of types that pool on some p it must hold that $\int_\tau u(a^R(t))f(t|\tau)dt - \max\limits_a \int_\tau u(a^R(t))f(t|\tau)dt \geq k^R\gamma'(1)$. If k^R is too high this inequality does not hold for any τ (since the left side of the inequality is bounded). Hence, if the reading cost is too high there is no equilibrium in which

[12] Indeed, the fact that it is better for all sender types makes the NRSE look susceptible to some refinements, such as the concept of announcement proofness introduced in Matthews et.al. (1991).

the receiver reviews reports with positive effort. A low k^R is therefore a necessary condition for a RSE to exist. The following result shows that it is also a sufficient condition

Proposition 4.6. *For k^R low enough there is a RSE.*

Proof. As in the non-reading SE, in the reading equilibrium type $t = 0$ must send $([0, 1], 1)$. Suppose $p(t) = 1$ for all $t \in T$. Then if k^R is sufficiently low, the receiver will examine p with effort $e = 1$. Further, let $m(t) = (l(t), 1)$ for all $t \in T$. Then $U^S(t) = u^S(a^R(t)) - (k^p\psi(-l(t)) + kC(1-l(t)))$. In a separating equilibrium we must find $l(t)$ such that for all $t, t' \in T$ such that $t' > t$ either $U^S(t) = U^S(t')$ or $t < l(t')$. But we know from Proposition 4.1 that such an $l(t)$ exists. To be explicit, let $\bar{t}^*(t) := \max\{\arg\max_{t'\in[0,t]} u^S(a^R(t')) - (k^p\psi(-t') + kC(1-t'))\}$ and let $\delta(t) := \{t' \in [0, t] : u^S(a^R(t)) - (k^p\psi(-t') + kC(1-t')) = u^S(a^R(\bar{t}^*(t))) - (k^p\psi(-\bar{t}^*(t)) + kC(1-\bar{t}^*(t)))\}$. Then $\delta(t)$ is similar to $\varepsilon(t)$ in Section 3, the only difference being that $\delta(t)$ is defined with a different cost function, that nevertheless satisfy the same properties as C. The arguments given in the proof of Proposition 4.1 thus applies here as well, so $m(t) = [\delta(t), 1]$ and $p = 1$ indeed is a reading SE if k^R is sufficiently small, under the additional condition that the receiver will not read any report (m, p') with $p' \neq 1$. $\square$

The proof of Proposition 4.6 shows that if the cost of examining a report is sufficiently low, it is always possible for the sender to pool on some appearance, and next separate using a combination of disclosure of information and costly signaling. While the appearance is the same for all types, so all reports look the same, the content of the reports differ across types. Once the receiver examines the reports, she is able to determine exactly which type sent the report. An equilibrium always exists due to the existence result in Section 3. The only difference here is that the sender incurs a combination of precision and manipulation costs, so the cost function is different (in fact, the sender indeed incurs higher costs here).

Hence, as long as the examination costs are not too high, perfect communication can occur in a way similar to the benchmark case studied in Section 3. A paradoxical aspect of the RSE specified in the proof of Proposition 4.6 is that all sender types except $t = 0$ manipulate the appearance of the report at a cost, with the only benefit that the receiver in this way will read their report. Moreover, they make their reports look less precise than what they are (perhaps by condensing the information, making it fit in a smaller number of pages). The problem that they face is that the receiver only examines reports with $p = 1$ and responds skeptically to any other report. This leads the sender to incur costs that seem unnecessary. The underlying dilemma arises since the receiver must read the reports for information to become verifiable and then to pooling on appearances is required. Manipulation is therefore necessary for separation. Hence, any separation arising from the receiver reading the reports requires that some types polish the appearances of the reports.

The receiver is worse off in the RSE than in the NRSE, since she obtains full information as before, but now at a positive cost. The sender incurs less costs than in the non-reading equilibrium, except in the special case in which $\bar{t}^*(1) = 0$ so the utilities across sender types is constant[13]. In this case, welfare is higher in the NRSE, since the expected utility of the sender is the same and the expected utility of the receiver is higher.

A conclusion is that perfect communication is possible even when the receiver has to make an effort at a cost in order to access the information contained in the report. RSE always exist if the intensity of the examination cost is not too high. This means means that the existence of a separating equilibrium is robust to the inclusion of small reading costs.

4.6 Concluding Remarks

In this paper it has been shown how the unraveling result of Milgrom (1981) generalizes when reporting of information increases continuously in the precision of the report. Contrary to what one might suspect, a separating equilibrium always exist, also for arbitrarily high reporting costs. Hence, communication can be perfect,

[13]The term $\bar{t}^*(1)$ is defined in the proof of Proposition 4.6.

even when communication costs are arbitrarily high. The intuition is that the costs work as a signaling device and a combination of disclosure of information and costly signaling can be used to accomplish full separation. When the costs of reporting become high the separating equilibrium looses uniqueness and other outcomes become plausible as well. Hence, when reporting costs are high, both situations in which the receiver obtains full information and situations in which she obtains less (or none) information can arise. This conclusion contrasts with the previous literature, such as Jovanovic (1982) and Verecchia (1983) which has argued that reporting costs are more detrimental for communication.

This paper has also studied the obstacles to communication that can arise when the receiver has to make an effort at a cost in order to assimilate the information contained in reports. It has been shown that there are both separating equilibria in which reports are not read and in which they are read. The former equilibria exist under a number of conditions on the cost functions. In principle, the precision and manipulation costs should not be too low. The latter exist as long as it is not too costly for the receiver to examine the reports. This means that unraveling is robust to costly examination of reports, in the sense that the set of separating equilibria is not emptied with small examination costs.

A promising direction for future research is to study more closely the role of the receiver in communication. For example, if the receiver must make an effort in order to understand the report, a relevant issue is the incentives of the sender to make it understandable. Another possibility is that the probability of understanding a report is related to its precision. In principle, a more precise report may be harder to understand, which would seem to reduce the incentives to produce precise reports. Finally, yet another possibility is to incorporate more "layers of information" into a report. For example, the receiver may obtain a first impression by a quick glance, a deeper understanding by browsing the report and finally a thorough understanding by reading it.

Bibliography

[1] Austen-Smith, D. (1994). Strategic Transmission of Costly Information. *Econometrica 62*, 955-963.

[2] Cheong I. and Kim J.-Y. (2004). Costly Information Disclosure in Oligopoly. *The Journal of Industrial Economics 51*, 121-132.

[3] Dewatripoint, M. and Tirole, J. (2005). Modes of Communication. *Journal of Political Economy 113*, 1217-1238.

[4] Eso, P. and Galambos A. (2007). Disagreement and Evidence Production in Pure Communication Games. Mimeo.

[5] Henry, E. (2009). Strategic Disclosure of Research Results: The Cost of Proving Your Honesty. *The Economic Journal 119,* 1036-1064.

[6] Jovanovic, B. (1982). Truthful Disclosure of Information. *The Bell Journal of Economics 13*, 36-44.

[7] Kartik, N. (2009). Strategic Communication with Lying Costs. *The Review of Economic Studies 76*, 1359-1395.

[8] Kartik, N., Ottaviani, M. and Squintani, F. (2006). Credulity, Lies and Costly Talk. *Journal of Economic Theory 134*, 93-116.

[9] Mathis, J. (2008). Full Revelation of Information in Sender-Receiver Games of Persuasion. *Journal of Economic Theory 143*, 571-584.

[10] Matthews, S., Okuno-Fujiwara, M. and Postlewaite, A. (1991). Refining Cheap Talk Equilibria. *Journal of Economic Theory 55*, 247-273.

[11] Milgrom, P. (1981). Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics 12*, 380-391.

[12] Milgrom, P. (2008). What the Seller Won't Tell You: Persuasion and Disclosure in Markets. *Journal of Economic Perspectives 22*, 115-131.

[13] Milgrom, P. and Roberts, J. (1986). Relying on the Information of Interested Parties. *RAND Journal of Economics 17*, 18-32.

[14] Seidmann, D. and Winter, E. (1997). Strategic Information Transmission with Verifiable Messages. *Econometrica 65*, 163-169.

[15] Verecchia, R. (1983). Discretionary Disclosure. *Journal of Accounting and Economics 5*, 179-194.